

Putting *oughts* together*

David Boylan
Texas Tech University

Submitted 2021-04-06 / First decision 2021-06-23 / Revision received 2021-10-17 /
Second decision 2022-01-09 / Revision received 2022-09-30 / Accepted 2022-11-27 /
Final version received 2023-10-01 / Published 2023-10-06 / Final typesetting 2024-
04-17

Abstract Consistent Agglomeration says that, when ϕ and ψ are consistent, \lceil ought ϕ \rceil and \lceil ought ψ \rceil entail \lceil ought $(\phi \wedge \psi)$ \rceil ; I argue this principle is valid for deontic, but not epistemic oughts. I argue no existing theory predicts these data and give a new semantics and pragmatics for *ought*: *ought* is an existential quantifier over the best partial answers to some background question; and presupposes that those best partial answers are pairwise consistent. In conjunction with a plausible assumption about the difference between deontic and epistemic orderings, this semantics validates Agglomeration for deontics but not epistemics.

Keywords: Agglomeration, weak necessity modals, ought, should, deontic, epistemic

The logic of *ought* is unruly, or seems so. Consider the principle *Consistent Agglomeration*:

Consistent Agglomeration: When ϕ and ψ are consistent:
 \lceil ought ϕ \rceil, \lceil ought ψ $\rceil \models \lceil$ ought $(\phi \wedge \psi)$ \rceil

(When there is no risk of confusion, I simply call it ‘Agglomeration’.) This is a multi-premise closure principle for *oughts*: when two consistent things

* I thank audiences at MIT, LENLS 15, the 2019 meeting of Pacific APA and the University of Texas, Austin and Fabrizio Cariani, Sam Carter, Kai von Fintel, Melissa Fusco, Benj Hellie, Jeff King, Justin Khoo, Matt Mandelkern, Shyam Nair, Milo Phillips-Brown and Stephen Yablo. Special thanks to Bob Stalnaker and Ginger Schultheis.

ought to be, so too ought their conjunction. The classic, necessity operator view of *ought* validates this principle. But as with much of standard deontic logic, various heterodox theorists argue Agglomeration succumbs to counterexamples.

Both traditions get something right and something wrong, I argue. The classic view is right that deontic *oughts* agglomerate. Nonetheless, Agglomeration is invalid, as the heterodoxy says — *epistemic* readings of *ought* do not agglomerate.

I give a new semantics for *ought*, where deontic and epistemic *oughts* genuinely have different logics. This arises from the different properties of deontic and epistemic orderings and how they interact with a novel *definedness constraint* for *ought*. In slightly more detail, I say that *ought* quantifies over the best answers to a *relevance question*, a question supplied by context to track the distinctions we take to be relevant. I add that *ought* requires the best propositions in context to be *pairwise consistent*. Finally I argue propositions can be epistemically but not deontically worse than all the relevant ways for them to come about. These pieces together are exactly what yields the difference in logic.

1 *Ought*, deontic and epistemic

In this section, I'll argue Agglomeration is valid for deontic but not epistemic *oughts*.

1.1 Flavours of *ought*

Readings of *ought* and *should* fall into two classes.¹ First, there are *deontic readings*. Following Charlow & Chrisman (2016), I take these to include readings concerning what we should do in light of *morality, law* or *practical reasoning*.² Take the sentence

- (1) John ought to/should be here by 10.

¹ Following the orthodoxy, I treat *ought* and *should* as essentially synonymous. See von Fintel & Iatridou's (2008) approach where the two words are “near-synonyms” and have the same semantics. Portner (2009) and Rubinstein (2012, 2014) also presuppose they are equivalent.

² Note I will use term ‘deontic’ more inclusively than the linguistics literature does.

Putting *oughts* together

This has a number of deontic readings: if John has made a promise to be here by 10, it is a moral *ought*; if John is a juror for a trial that starts at ten, it might be a legal *ought*; if we are at the airport and John is on his way to catch a flight, it might express a prudential *ought*.

But *ought* also has a non-deontic reading. (1) might also be heard to be true because the bus usually takes 30 minutes to reach our stop and it is now 9.30. This reading does not concern whether it would be *good* or *required* for the bus to arrive in 10 minutes. It expresses an *expectation* about when the bus will arrive. Call this an *epistemic* reading of *ought*.

Other things being equal, we should seek a unified account of the epistemic and deontic *ought*. When a single modal word expresses a variety of modal flavours, both in English and other languages, this is unlikely to be a chance ambiguity.³ Better to follow Kratzer's (1977) example: give a unique semantic entry and account for the variation in flavour with different contextually supplied parameters.

1.2 Epistemic Agglomeration

Comparing the *logic* of epistemic and deontic *oughts* pushes in a different direction. Epistemic Agglomeration can fail when we expect each of a set of propositions to be true, but not their conjunction.⁴

Take the following example:⁵

The Office. 26 workers, Alice, Bob, Carol, ..., and Zadie, work on separate floors of an office building. On average they take a sick day once a month: there is a regular, low-level circulation

³ See von Stechow & Iatridou 2008 for a survey of the cross-linguistic data here.

⁴ This is the structure of Makinson's (1965) *preface paradox*.

⁵ The editor points out that we often seem to rely on epistemic Agglomeration, in straightforward cases. They give the following example:

- (i) A. How long do you think it's going to take us to get to the movie theater?
- B. Well, the bus should be here in about ten minutes, the ride should take about 15 and then it should only take another 5 mins to walk to the theater, so we should be there in half an hour.

I agree with their judgement in this case. The issue is that epistemic Agglomeration failures tend to require many conjuncts. Even in *The Office*, we can infer that Alice and Billy should be in. But as we add more conjuncts, and so increase the risk of a proposition being false, we become less willing to agglomerate. (This is also a feature of the original preface paradox.)

of various colds and viruses through the building, so that, statistically, it is rare that all of the 26 workers are in the building on a given day.

The following are all true here:

- (2) Alice should be in the office today.
- (3) Bob should be in the office today.
- (4) Carol should be in the office today.
- ...
- (5) Zadie should be in the office today.

But, when *ought* takes wide scope, it does not follow that:

- (6) Everyone should be in the office today.

Agglomeration predicts otherwise: If (2) — (5) are true, then (6) should be too.

Here is one more example, adapted from [Carter & Hawthorne 2021](#):

Life Expectancy. In the US it is statistically unlikely for anyone now a teenager to die before the age of 30. But it is also statistically practically guaranteed that *some* teenager will die before the age of 30 each year.

Of each individual teenager in the US, it is true to say:

- (7) They should not die before they are 30.
- (8) They should at least live past the age of 30.

But it is not true to say:

- (9) Every teenager should live past 30.
- (10) It should be that every teenager lives past 30.

1.3 Deontic Agglomeration

On the other hand, various considerations suggest deontic Agglomeration is valid.

Putting *oughts* together

First, Agglomeration is hard to resist in straightforward cases. Suppose I ought to help Alice and I ought to help Billy. It is hard *not* to conclude I ought to help Alice and help Billy — it is eminently plausible that if there are two things, each of which I ought to do, then I ought to do *both*. And we can keep agglomerating, as we add further obligations. If I should help Alice, Billy *and* Carol, or Alice, Billy, Carol *and* Daniel, then in each case I ought to help all of them. And so on.

Second, discourses that violate Agglomeration tend to sound contradictory out of the blue. Consider:

- (11) a. You may not go to the movies and to Johnny's house.
- b. But you should go to the movies and you should go to Johnny's house.

This would violate Agglomeration: it's generally accepted that *ought* entails *may*; so (11a) will entail that it is not the case that you ought to go the movies and to Johnny's house. For this very reason, (11a) seems inconsistent with (11b).

Third, the contrapositive of deontic Agglomeration seems valid. Consider the following case:

Painkillers. Alice has a headache and there is both Advil and Aleve in the house. She always forgets which painkillers can be taken together and which painkillers she finds individually most effective. But Billy often remembers such things.

Imagine the following conversation ensues:

- (12) Alice: My headache is truly terrible. Should I take Advil *and* Aleve?
- Billy: Absolutely not; it's not safe to take both.
- Alice: Well then, should I take Advil? Should I take Aleve?

If Billy knows that Advil is more effective for Alice, he could simply say:

- (13) Yes, you should take the Advil.

In that case, if he replies to Alice's final question, he is clearly committed to answering:

- (14) You should not take the Aleve.

Alternatively, suppose Billy does not remember whether either is more effective. He could say:

- (15) I don't remember if there is a particular one you should take. But if there is, then you should not take the other one.

Both patterns are explained by Agglomeration. Given deontic Agglomeration, if it is *false* that Alice ought to take Advil and Aleve, then there is at most one painkiller she should take; if there is a particular painkiller she should take, it follows that it cannot be true that she should take the other. This predicts why, if Billy first replies with (13), he is committed to (14). It explains also why Billy is in a position to assert the second conjunct of (15).

I submit that Agglomeration is valid for the deontic *ought*. But there is also a further principle that, while not strictly speaking an entailment of Agglomeration, is naturally expected to be valid too. To work up to the principle, consider the following example:

Dessert. There are three dessert options: cannoli, cheesecake, and apple pie. Pie and cannoli are tastiest. I can order as many dishes as I like, but I will definitely feel ill if I have more than one.

Here the conjunctive *ought* claim below is clearly false:

- (16) I ought to have pie and cannoli.

After all, having both will make me ill. Agglomeration then tells us that at most one of the following claims can be true:

- (17) I ought to have pie.
(18) I ought to have cannoli.

But in fact, both seem false: if they are equally good, how could one be true and not the other? It seems that the strongest assertable *ought* claim is disjunctive:

- (19) I ought to have pie or cannoli.

If the options are equally good but I should not take both, then all that I *should* do is bring about the disjunction.

Putting *oughts* together

Examples like this motivate the principle that multiple incompatible but deontically best options make for disjunctive deontic *oughts*:

Indifference: When ϕ_1, \dots, ϕ_n are each (intuitively) deontically best, on the deontic reading $\lceil \text{ought } (\phi_1 \vee \dots \vee \phi_n) \rceil$ is true and $\lceil \text{ought } \phi_i \rceil$ is false for $1 \leq i \leq n$.

While Indifference is not entailed by deontic Agglomeration, they make for a very natural package. In cases of multiple best options, it is hard to see what other advice an Agglomeration defender could give: the best we can say is that one of the best options should be brought about, but that it is left open which particular one; any other advice seems arbitrary, if the options are truly equally good. So in addition to deontic Agglomeration, I will aim to predict Indifference too.

2 Defending the data

In this section, I deepen the case for the generalisations from the previous section. First, I consider various purported counterexamples to deontic Agglomeration. I argue these examples are really due to the context-sensitivity of the deontic *ought*: once we make distinctions familiar from the literature on deontic logic, the deontic counterexamples are seen to rest on an equivocation, an equivocation that has no parallel in the epistemic case. I then consider an argument in favour of epistemic Agglomeration: certain cases structurally similar to The Office fail to be counterexamples to Agglomeration; and the data observed in these cases would be well explained if epistemic Agglomeration were valid. I instead argue that these cases should be explained by appeal to the fact that high probability does not necessarily suffice for an epistemic *ought*, a claim that has been independently motivated by Copley (2006) and Yalcin (2016).

2.1 Chariots

Jackson (1985) claims the following example fells deontic Agglomeration:

Chariots. Attila and Genghis are driving their chariots towards each other. If neither swerves, there will be a collision; if both swerve, there will be a worse collision . . . but if one swerves and the other does not, there will be no collision. Moreover if

one swerves, the other will not because neither wants a collision. Unfortunately, it is also true to an even greater extent that neither wants to be 'chicken'; as a result what actually happens is that neither swerves and there is a collision.

Jackson says the following are true here:

- (20) Attila ought to swerve.
- (21) Genghis ought to swerve.

But the following is clearly *false*:

- (22) Attila and Genghis both ought to swerve.

This would be contrary to Agglomeration.

The counterexample presupposes that *ought* is univocal throughout. But is it? Philosophers distinguish between two families of readings of the deontic *ought*, the *ought to do* and the *ought to be*.⁶ The *ought to do* states something that is required of or recommended to a particular agent; and the *ought to be* states how the world would be if things went best.

These sometimes come apart. Suppose Alice needs help packing for a move and of her friends only Billy is willing to help. As a result, the packing will take all day. The following might be true here:

- (23) Billy should spend his entire day helping Alice.

But the denial of the corresponding ought-to-be *also* seems true:

- (24) It ought not be that Billy spend his entire day helping Alice.

Things would be better and fairer if Alice's other friends came to help too, speeding up the packing.

Once distinguished, the counterexample fails. The true readings of the premises are most naturally understood as involving the *ought to do*. But the *ought to do* is itself a *family* of readings, not an individual reading: there is an *ought to do* reading for each individual agent. For (20) and (21) say the same thing as:

- (25) Attila should make it the case that Attila swerves.

⁶ See for instance [Schroeder 2010](#).

Putting *oughts* together

(26) Genghis should make it the case that Genghis swerves.

But when we paraphrase (21) in terms of *Attila's ought to do* it is simply false:

(27) Attila ought to make it the case that Genghis swerves.

Likewise, when we paraphrase (20) in terms of *Genghis's ought to do*:

(28) Genghis ought to make it the case that Attila swerves.

This suggests an equivocation in the premises, even when they are read with the *ought to do*. There can be many *ought to dos*, even in a given case. There is what *Attila* ought to do and what *Genghis* ought to do. But on neither reading are *both* of the premises true.

When we elicit the *ought to be* reading, the premises are no longer clearly true:

(29) It ought to be that Attila swerves.

(30) It ought to be that Genghis swerves

Because while (exactly) one of them should swerve, it does not matter which: things go best if (just) Attila swerves or if (just) Genghis swerves. So, on the *ought to be*, the counterexample falters too.

2.2 Layover

An anonymous reviewer suggests a different counterexample:

The Layover. I have a twelve hour layover at Paris and will use it to see some of the city. I ask my friend Alice who has been to Paris many times what I should see there. She replies:

- (31)
- a. You should see Centre Pompidou, because the art there is incredible.
 - b. And you should see the Tour Eiffel because it's the icon of the city.
 - c. And you should see the Louvre because it's the most famous museum in the world.
 - d. But you really shouldn't do *all* of these, as you need to be back at the airport at least two hours before your flight.

Again I think there are two different kinds of *oughts* in play. (31a) – (31c) talk about what you *prima facie* ought to do; but (31d) talks about what you ought to do *all things considered*.

Let's first consider what this distinction amounts to.⁷ Take a claim like:

(32) If you're in the neighbourhood, you should drop by.

I might truly say this to a friend. But clearly, I am assuming various things in the background: for one thing, I am assuming that if my friend is in the neighbourhood, it won't be 2am. Roughly speaking, I am assuming things are normal: *everything else being equal*, they should drop by, if they are in the neighbourhood. This is the *prima facie* *ought*.

But of course, everything else is not always equal. Suppose my friend ends up in my neighbourhood at 2am and they call me. I can perfectly consistently say:

(33) You should not drop by; it is too late.

Here I am holding fixed *all* the facts as they actually are; I am not assuming everything else is equal. This is the *all things considered* *ought*.

Why should we think that The Layover has anything to do with this distinction? Firstly, because once (31d) is uttered, we can easily elicit the intuition that one of the premises must in fact be false. Given all of what Alice said, I would be within my rights to reply:

(34) a. Ok, so there must be at least one that I should not try to see.
b. Which ones should I *actually* try to see?

And we could imagine Alice replying in several different ways:

(35) You should see *just* the Louvre; forget about the others.

(36) You should skip Centre Pompidou.

(37) I don't know which ones you should actually see; it's too hard to choose.

What she *can't* do is insist that, by saying (31a) – (31c) she has already answered the question; my question in (34b) is *not* redundant, (31a) – (31c) notwithstanding. This is indicative of a context-shift, one well explained by

⁷ This distinction ultimately traces back to Ross 1930.

Putting *oughts* together

a shift from the *prima facie* to the all things considered *ought*: the *ought* in my question is not the same as that in (31a) – (31c)

But why think that the context-shift specifically happens between the premises and the conclusion? Because, once (31d) is truly spoken, it is legitimate for me to *infer* that we are now in a context where one of the premises is false. Indeed, this is exactly what I am doing, when I utter (34a). That I can say (34a) is hard to explain *without* Agglomeration. If *oughts* do not Agglomerate, I have no right to this inference. And the inference in (34a) is what prompts my question in (34b) in the first place.

As a final piece of evidence, observe we can elicit a true *prima facie* reading of the agglomerated premises, when we add into the mix a destination that we should *prima facie* avoid:

- (38) You should see Centre Pompidou (because of the fantastic art), the Tour Eiffel (because it's an icon) and the Louvre (because it's the most famous museum in the world). (You shouldn't bother with the Arc de Triomphe — it's a real let-down.)

The agglomerated *prima facie* reading should be easier to access, when we make clear there are some attractions that should be avoided (even *prima facie*). And indeed, I think it is.

I submit then that, when we look at the bigger picture, *The Layover* is also best explained by an equivocation, this time between the *prima facie* and all things considered *ought*.

2.3 Context-shifting in *The Office*?

By distinguishing different readings of *ought*, we defused the challenge of *Chariots* and *The Layover*. But could the judgements in *The Office* also be due to a context-shift?

Unlike in *Chariots*, this move overgenerates here. Suppose that some premises are evaluated in a different context from the conclusion. If Agglomeration were valid, one of the premises would be *false* in the context of the conclusion. So in that context at least one of the *negations* of the premises in *The Office* should be true:

- (39) It's not true that Alice should be in the office today.
(40) It's not true that Bob should be in the office today.

(41) It's not true that Carol should be in the office today.

...

(42) It's not true that Zadie should be in the office today.

I don't hear any true reading of these claims in The Office. And if they existed, we should be particularly primed to hear them after denying the conclusion of the inference: in that context, at least one will be true; and considerations of charity would suggest that we will *remain* in such a context, if it makes true the sentences under consideration. What's more, further pressure does not elicit the relevant readings. It still sounds false to say:

(43) It's not true that *Alice* should be in the office today. (After all, not everyone will be in.)

Not so in Chariots: there we can easily hear as true the sentence:

(44) It's not true that *Attila* ought to swerve. (After all, *Genghis* could swerve.)

But perhaps the missing readings are elusive in the sense of Lewis 1996. Perhaps when I consider an individual premise, I am moved into a context where it is true, and where a premise not under consideration is false. The false readings would exist, but never where we are looking.⁸ This would also explain why (6), the conclusion of the inference, has no true reading in The Office: on the elusive strategy, there is no one context that makes *all* the premises true.

These elusive false readings should still be indirectly accessible with quantifiers. The following wide-scope *universal* claims would be univocally false:

(45) Everyone is such that they should be in the office today.

(46) Everyone is someone who should be in the office today.

After all, the universal quantifier has an (elusive) counterexample in *every* context. And it should be univocally *true* to say:

(47) Not everyone is someone who should be in the office today.

⁸ Hawthorne (2002) explores this response to the preface paradox for the case of 'knows'.

Putting *oughts* together

Finally, upon realisation that somebody is very likely to be absent, it should make sense to ask:

(48) So who then don't you think should be in the office today?

These predictions are incorrect. (45) and (46) are in fact true; (47) is false; and (48) rests on a false assumption about the case, namely that (47) is true. Notice as well that (45) and (46) should be false in *every* context, even for this kind of contextualist; and (47) should be true in every context. For (6) is after all univocally false: and so in *every* context, there is an elusive counterexample to claims like (45) and (46).

2.4 Absent counterexamples?

Finally, a different kind of worry about the epistemic counterexamples: they are fewer than one might have expected; we do not get counterexamples in just any case where probability fails to agglomerate.

An anonymous reviewer draws attention to a case analogous to The Office:

High School Musical. 26 students are in the last round of auditions for a high school musical. Each student has verbally promised that, if selected, (s)he will be available for rehearsals starting at 7am sharp every day from tomorrow through opening night. Uncomfortably, there are 25 roles! The one student who will be cut from the roster will be notified by email at midnight tonight. It would be painfully awkward for everyone involved, including the director, if the student who is cut showed up tomorrow for rehearsal only to be sent home.

On its epistemic reading, the conclusion of a natural Agglomeration inference is clearly false here.

(49) All the students should be in the gym tomorrow at 7am.

But, the reviewer notes, correctly in my view, that the analogous premises sound false in this case:

(50) Anna should be in the gym tomorrow at 7am.

(51) Bob should be in the gym tomorrow at 7am.

...

(52) Zadie should be in the gym tomorrow at 7am.

This might be taken as some evidence in favour of epistemic Agglomeration. We saw it is clearly false that all the students should be in the gym. If Agglomeration were valid, then at least one of the premises must fail. But our evidence concerning each student is on a par: for instance, we have no more reason to think Anna will be in than Bob. So it is natural to think if any premise fails, they must *all* fail. And that is exactly what we see. If Agglomeration is not valid, then why else would the premises be false?⁹

I think the best overall explanation here is different. The falsity of the premises in High School Musical is explained not by Agglomeration, but rather by the fact that high probability alone does not tend to suffice for the truth of an epistemic *ought*. This observation is due to Copley (2006) and Yalcin (2016). For instance, despite the low odds of winning, it sounds off to say things like:

(53) I shouldn't win the lottery.

(54) I ought to lose the lottery.

I think that this explains why we are not inclined to accept the premises in High School Musical. The only reason we would think it likely that, say, Anna has been selected for a part in the musical is the simple fact that only one student will be cut. But this kind of evidence does not suffice for an epistemic *ought*.

Now I don't mean at this point to take any further stand on the relationship between probability and the epistemic *ought*. It might be that, as Yalcin (2016) proposes, the epistemic *ought* really concerns the degree of normality of a proposition; or it might be that it concerns some other relation, call it *expectation*, of which high probability is a necessary but not sufficient con-

⁹ An anonymous reviewer notes that, like in Chariots, we are inclined to judge the following true here:

- (i) It's not the case that Anna should be at the gym tomorrow at 7am. After all, she could be the one who got cut!

This seems to me a good way to bolster the case that High School Musical is not an Agglomeration failure.

Putting *oughts* together

dition. But notice either notion would substantiate my claim here. In High School Musical it is not natural to say either of:

(55) I expect Anna to be in the gym tomorrow at 7am.

(56) It would be normal for Anna to be in the gym tomorrow at 7am.

Why, on the other hand, are the epistemic *ought* claims true in The Office? I think it is because here the probabilities do not simply result from counting possibilities. Rather, probability is a good indicator of what is normal and what we should expect. For instance, given the information about absences and illness in the office, it would be very natural to say either of:

(57) I expect Alice to be in the office today.

(58) It would be normal for Alice to be in the office today.

So, in this case, probability and normality *do* align. And similarly, here our expectations *do* seem to be based on the statistics. For this reason, the relative probabilities are a good guide to what epistemic *ought* claims are true. So, while I do not in general assume that probability suffices for an epistemic *ought*, I do assume that statistical information is a guide to whatever ordering over possibilities is relevant in cases like The Office. Given that it is not obvious what the Agglomeration defender should say about The Office, I think denying epistemic Agglomeration still provides the best overall explanation here.

3 The predicament

Agglomeration fails for epistemics, but not deontics. I show that the main views in the literature fail to explain this.

3.1 The classic semantics

On the classic semantics, *ought* is a necessity modal. To fix ideas, take von Fintel's (2012) Kratzerian statement of it:

(59) $\llbracket \text{ought } \phi \rrbracket^w = 1$ iff $\forall w' \in \text{BEST}(w): \llbracket \phi \rrbracket^{w'} = 1$

That is, ‘ought ϕ ’ is true iff ϕ is true throughout some set of best worlds. What makes a world best does not matter: (59) alone determines the logic of *ought*.

The classic semantics predicts Agglomeration is simply valid. If ‘ought ϕ ’ and ‘ought ψ ’ are true, then all the best worlds are ones where ϕ is true and where ψ is true. But then all the best worlds must be ones where ‘ $\phi \wedge \psi$ ’ is true. So ‘ought $(\phi \wedge \psi)$ ’ is true. This reasoning assumes nothing about the flavour of the *ought*. Any counterexample to Agglomeration undermines the classic view, when understood as a unified semantics.

The classic semantics has limited options. It might pursue an ambiguity treatment for *ought*. But this sits badly with the known fact that *ought* and modal vocabulary more generally express a variety of different flavours across languages. It might deny there really is a difference in logic across flavours. But that sits badly with the data from Section 2.3. If Agglomeration really does fail for *ought*, the classic semantics is off the table.

3.2 Contrastivism

Existing *contrastivist* accounts of *ought* fare no better.¹⁰ Contrastivists think that ‘ought ϕ ’ is true iff ϕ is *better* than its alternatives.¹¹ Slightly more precisely, where $ALT(p)$ is the set of p ’s alternatives and $\prec_{w,f,g}$ some ordering over propositions:

$$(60) \quad \llbracket \text{ought } \phi \rrbracket^{w,f,g} = 1 \text{ iff for every } q \in ALT(\llbracket \phi \rrbracket^{f,g}): \llbracket \phi \rrbracket^{f,g} \prec_{w,f,g} q$$

Even closely related propositions have different alternatives. In particular, the alternatives for p and for q might be different from those for $p \wedge q$.¹²

Many contrastivists also want to invalidate Agglomeration.¹³ They do so when p is better than its alternatives, q is better than its alternatives but $p \wedge q$ is not better than its alternatives. Most contrastivist views invalidate *deontic* Agglomeration. On most natural ways of thinking about alternatives,

¹⁰ Section 9 explores a new form of contrastivism which borrows some of the machinery of my view.

¹¹ In addition to Jackson 1985, see also Goble 1996, Finlay 2009, Lassiter 2011 and Snedegar 2012.

¹² A point on notation: I use lower-case Roman letters as variables with propositions as values; Greek variables take sentences as their values; and $\llbracket \phi \rrbracket^{f,g}$ is the proposition expressed by ϕ , relative to $\langle f, g \rangle$.

¹³ Many, but not all. Cariani (2013b,a, 2016b) defends a contrastivist semantics which validates Agglomeration in full generality.

Putting *oughts* together

p and q can both be *deontically* better than their alternatives, while $p \wedge q$ is not.

For example, Jackson says that $ALT(p)$ is just $\{\neg p\}$; and that p is better than q if the closest world where p is true is better than that where q is true.¹⁴ This makes Chariots a counterexample to Agglomeration. The premises of the inference are true: the closest world where Genghis swerves is one where Attila remains on course; it must then be better than the closest world where he doesn't and both remain on course. Likewise for the closest worlds where Attila swerves and where he does not. But the conclusion is false: neither swerving leads to the worst possible outcome and so the closest world where both swerve is worse than that where at least one does not.

This is of course by design. Chariots is meant to motivate Jackson's view. But this is not what the data support, I argued. Only the epistemic case gives robust counterexamples to Agglomeration.

3.3 The conflict account

The *conflict account* of *ought*, defended in van Fraassen 1973, von Fintel 2012, Horty 2012 and Swanson 2016, says that 'ought ϕ ' is true, roughly, just in case some maximally consistent subset of a set of best propositions entails ϕ .¹⁵ While it inspires elements of my own view, the conflict view has two important shortcomings: it *undergenerates* epistemic Agglomeration failures; and it incorrectly predicts epistemic Agglomeration failures involve dilemmas.

Again I will follow von Fintel's (2012) presentation. Say that $g(w)$ returns a set of best propositions at w ; and $f(w)$ is the background information in w . Now define $D(f(w), g(w))$ to be the set of maximal contextually-consistent subsets of $g(w)$.¹⁶ Then *ought* has the following entry:

$$(61) \quad \llbracket \text{ought } \phi \rrbracket^{w,f,g} = 1 \text{ iff for some } S \in D(f(w), g(w)) \\ \text{for all } w' \in \bigcap S: \llbracket \phi \rrbracket^{w',f,g} = 1.$$

¹⁴ See also Cariani 2016a, which gives a clear account of why, on various natural ways of thinking about alternatives, Agglomeration fails on the contrastivist accounts of Finlay 2009 and Lassiter 2011, among others.

¹⁵ Only Swanson explicitly defends the conflict account for epistemics. But, unless they reject a unified semantics, these authors are committed to applying it to epistemics.

¹⁶ In other words, the set of subsets S such that $S \cap f(w)$ is consistent and there is no S' such that $S \subset S' \subseteq g(w)$ and $S' \cap f(w)$ is consistent.

The conflict account primarily aims to invalidate Agglomeration failure for *inconsistent oughts*: 'ought ϕ ' and 'ought $\neg\phi$ ' may be consistent on the conflict account. But it also invalidates Consistent Agglomeration.¹⁷

To get the problems going, let's first see how the conflict account can be extended to epistemics. I assume that *epistemically* best propositions are ones that pass some threshold for likeliness or normality. In *The Office*, it is natural to assume these include *Alice is in*, *Billy is in*, ... , *Zadie is in* and *Not everyone is in*. This set is inconsistent, but it will have maximal consistent subsets like:

- (62) {*Billy is in*, *Carol is in*, *Daniel is in*, ..., *Zadie is in*, *not everyone is in*}
- (63) {*Alice is in*, *Carol is in*, *Daniel is in*, ..., *Zadie is in*, *not everyone is in*}
- (64) {*Alice is in*, *Billy is in*, *Daniel is in*, ..., *Zadie is in*, *not everyone is in*}
- (65) {*Alice is in*, *Billy is in*, *Carol is in*, *Daniel is in*, ..., *Zadie is in*}

The conflict account then predicts the premises from *The Office*, repeated below, are true.

- (2) Alice should be in the office today.
- (3) Bob should be in the office today.
- (4) Carol should be in the office today.
- ...
- (5) Zadie should be in the office today.

But the conflict account does *not* predict that the conclusion of the inference, repeated below, is false. It says it is univocally true:

- (6) Everyone should be in the office today.

After all, there *is* a maximal consistent subset which entails everyone is in, namely (65). The conflict account fails to predict that (2) - (6) is a counterexample to Agglomeration.

The second problem is that *The Office* is predicted to contain *epistemic dilemmas*: for some ϕ , both 'ought ϕ ' and 'ought $\neg\phi$ ' are true on their epistemic readings. Both of the following are predicted to be true:

¹⁷ **Proof:** Suppose $f(w) = W$, $g(w) = \{p, q, r\}$, $q \cap r \neq \emptyset$ and but $p \cap q \cap r = \emptyset$. Then $D(f(w), g(w)) = \{\cap\{p, q\}, \cap\{p, r\}\}$. Here $\llbracket \text{ought } q \rrbracket^{c, w, f, g} = \llbracket \text{ought } r \rrbracket^{c, w, f, g} = 1$, but $\llbracket \text{ought } (q \wedge r) \rrbracket^{c, w, f, g} = 0$

Putting *oughts* together

(2) Alice should be in the office today.

(66) Alice should *not* be in the office today.

We already saw why (2) is true. (62) above witnesses the truth of (66): given the background information, that set entails Alice is absent. So the conflict account predicts the following sentence is true:

(67) Alice should be in the office today and Alice should not be in the office today.

In fact, each premise generates a dilemma, because for each worker there is a maximal consistent set that (contextually entails) they are absent today. But The Office involves no dilemma. (66) is false; and (67) is simply incoherent here.

An anonymous reviewer notes that further resources from Horty 2012's conflict account might solve the problem. What I have called best propositions Horty thinks of as *defaults*. And on his view, some defaults can supersede others: the default *penguins don't fly* might be stronger than the default *birds fly*; and when defaults conflict, given the background information, we prioritise the stronger default. This requires a more complex algorithm to determine $D(f(w), g(w))$: we do not consider all the maximal consistent sets, but only the ones containing the strongest defaults.¹⁸ Now apply this to The Office: perhaps *not everyone is in* should be ranked higher than any of the other best propositions like *Alice is in*. This would have the result that (65) is not in $D(f(w), g(w))$ and so (6) would not be true.

There are two things to say here. The first is that the problem of dilemmas still applies to the more complex view; and I'm inclined to think that this is really the more serious problem. But second, I do not think that the proposed ordering is plausible on every construal of the case. For instance, suppose that on average one day a month everyone is in the office; thus, everyone is present just about as often as each particular person is absent. Here it's not obvious that *not everyone is in* should have a higher priority than, say, *Alice is in*: it's no more probable or normal. And for similar reasons it's not obvious that you are more committed to one or the other: if you were to learn that *exactly one* of the two propositions *not everyone is in* and *Alice is in* were true, you would not necessarily conclude that Alice is not in; especially so,

¹⁸ See Horty 2012 for the full statement of the algorithm.

when we are explicit that on average everyone is in the office about once a month.

4 The semantics

On my view, the difference between deontics and epistemics arises from how differing properties of their orderings interact with a consistency constraint. To fully flesh out this idea, I give both a novel semantics and pragmatics for *ought*.

In this section, I start with the semantics: I give the basic semantic entry and then enrich it in two steps, adding *question-sensitivity* and then a *pairwise definedness constraint*. The following sections discuss pragmatics: I introduce and motivate a special kind of contextually supplied question, a *relevance question*; and I motivate the crucial distinction between deontic and epistemic orderings.

4.1 The first pass

The basic idea behind the semantics is simple: ‘ought ϕ ’ is true iff ϕ is entailed by *some best proposition*. To state this more precisely, add a modal base, f , to the index; $f(w)$ is a set of worlds representing the background information in w .¹⁹ Add also an ordering function g ; this takes a world and a modal base to an ordering over propositions. (I write its output as $\prec_{w,f,g}$.) Now define the set of the best subsets of our information in w :

$$(68) \quad PBEST(w, f, g) = \{p \subseteq f(w) : p \neq \emptyset \text{ and } \neg \exists q \subseteq f(w) : q \prec_{w,f,g} p\}$$

I assume g outputs orderings that obey the Limit Assumption: we do not get infinite chains of increasingly good propositions.²⁰ This ensures that $PBEST$ is non-empty. I then say that *ought* quantifies over this set:²¹

$$(69) \quad \llbracket \text{ought } \phi \rrbracket^{w,f,g} = 1 \text{ iff } \exists p \in PBEST(w, f, g) : \forall w' \in p : \llbracket \phi \rrbracket^{w',f,g} = 1$$

What particular orderings does g output? A large and, to my mind, convincing literature argues a semantics for *ought* should aim for deontic neu-

¹⁹ Following Kratzer (1977), I allow modal bases to be either epistemic or circumstantial.

²⁰ More precisely: for every w, f, g , and for any set of propositions S there is some $p \in S$ such that $p \neq \emptyset$ and for no $p' \in S$ $p' \prec_{w,f,g} p$.

²¹ Moss (2015), Mandelkern, Schultheis & Boylan (2017) and Khoo (2021), among others, also propose quantifying over propositions.

trality: the lexical entry for *ought* should not be incompatible with plausible theories of moral value or decision theory. Best just to say some contextually relevant source of value supplies an ordering to the semantics.²² Our judgements of betterness in context will track this ordering.

What about epistemic orderings? Some suggest epistemic *oughts* track whether or not the *probability* of the prejacent passes a contextually supplied threshold. Others say that epistemic *oughts* track the relative *normality* of the prejacent: we would have $p \lesssim_{w,f,g} q$ just in case p obtains in more normal situations than q . In Section 2, we saw some reasons to be skeptical of the pure probability option: high probability alone does not in general suffice for an epistemic *ought*. Either option is compatible with my theory; but, as I have flagged, in various cases I will assume that the ordering for an epistemic *ought* shares important *structural* features of a probability ordering.

It is easy to see how Agglomeration can fail on my account. Let $[\phi]^{f,g}$ and $[\psi]^{f,g}$ be best and both better than any proposition entailing $[\phi \wedge \psi]^{f,g}$. In such a case ‘ought ϕ ’ and ‘ought ψ ’ will be true; ‘ought $(\phi \wedge \psi)$ ’ will not. But we need Agglomeration to fail only for epistemics, not deontics. My approach is to put further constraints on what orderings can be supplied to the semantics.

4.2 Questions

The set of propositions is very large indeed. Which ones get ordered in the semantics?

I say only answers to a contextually supplied question get ordered. (More in Section 5 on what particular question that is.) To spell this out, let’s assume a question is a set of propositions that partitions the background information into the complete answers to the question.²³ For instance the question *Who out of Alice and Billy had ice cream?* is a partition we could draw as follows:

²² See, among others, Carr 2015, Charlow 2016 and Cariani 2016b.

²³ Here I use Karttunen’s (1977) account of questions to model an issue relevant in the context. I take no stand on the semantics for interrogatives.

Just Alice	Just Billy
Both	Neither

A *complete answer* to Q is just one of its elements. A *partial answer* to Q is the union of some elements of Q . So, for instance, among the partial answers to the question *Who out of Alice and Billy had ice cream?* will be the proposition *either just Alice or both Alice and Billy had ice cream*.

I let take a question to supply the set of propositions that get ordered. Where Q is the relevant question, g orders just the partial answers to Q . This gives us a fairly sensible answer to our opening question: the final ordering only cares about propositions distinguished by the relevant question. A proposition does not get ordered if it cross-cuts distinctions made by the question.

To spell out this more precisely, we let ordering functions also take *questions* as arguments: g is now a function from worlds, modal bases and *questions* to an ordering $\prec_{w,f,g,Q}$. Say that $Q|S$ is the partition imposed by Q restricted to S ; that is,

$$Q|S = \{p : \exists q \in Q \text{ and } p = q \cap S\}.$$

Then define the *question-sensitive* ordering functions as follows:

Question-sensitivity: g is a *question-sensitive* ordering function iff $\prec_{w,f,g,Q}$ orders only complete and partial answers to $Q|f(w)$.

I stipulate that our semantics will draw from the set of question-sensitive ordering functions, rather than the set of ordering functions more generally. We then redefine *PBEST* with this in mind:

$$(70) \quad PBEST(w, f, g, Q) = \{p \subseteq f(w) : p \neq \emptyset \text{ and } \neg \exists q : q \prec_{w,f,g,Q} p\}$$

Question-sensitivity gives us a principled account of what propositions get ordered; and later will allow us to state intuitive constraints on orderings.

Putting *oughts* together

4.3 A consistency definedness condition

I propose that *ought* requires the best partial answers to be *pairwise contextually consistent*.

I adopt this as a definedness constraint: given a particular ordering, $\lceil \text{ought } \phi \rceil$ has a truth-value only if every pair of best propositions are consistent, given the background information:

Consistency Constraint: $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q}$ is defined only if for all p and q in $PBEST(w, f, g, Q)$, $(p \cap q) \cap f(w) \neq \emptyset$.

Note that this does not require that $p \cap q$ actually be a member of $PBEST(w, f, g)$.²⁴

It is important that we use *pairwise*, rather than overall, consistency. The Office illustrates why. There, on the natural way of construing the situation, the entire set of best propositions is

{Alice is in work today, Billy is in work today, ..., Not everyone is in}

This set is pairwise consistent: any two people could be in together; and any particular person could be in, even if somebody is absent. But it is not overall contextually consistent: somebody is out of work just in case one of Alice, Billy, Carol, ... and Zadio is absent.

²⁴ Ultimately I suspect this constraint needs to be strengthened, so not only must the best propositions be consistent, but also the propositions at each level below the very best too. Where $PBEST_n$ is the set of the n th best propositions, the condition would say:

Strong Consistency: $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q}$ is defined only if for all p and q in $PBEST_n(w, f, g, Q)$, $p \cap q$ is consistent with $f(w)$, for all n such that $PBEST_n(w, f, g, Q)$ is defined.

While I will continue to use the weaker constraint, I think there are at least two reasons to ultimately prefer Strong Consistency. First, Indifference should generalise to conditional consequents. As things stand, when the best propositions are consistent and the second best propositions are not, the following will be undefined:

- (i) If we don't perform a best option, we should perform a second best option.

While accommodation could account for this, I prefer to assume Strong Consistency Constraint: this will ensure from the get go that we have a question that makes both the conditional and unconditional *oughts* defined.

Second, as I discuss in Section 9, the stronger constraint is better able to predict natural values for the relevance question parameter in context.

What kind of definedness condition is this? I will think of it along the lines of a definedness condition on a pronoun.²⁵ Pronouns can carry gender, person or number features. Following Heim & Kratzer (1998), semanticists tend to model these as definedness constraints: the extension of a pronoun like ‘she’ is only defined on a given variable assignment if the extension is female. On the final analysis, items like modal bases and ordering functions will also be supplied by variable assignments; so this kind of definedness seems appropriate for constraining orderings.

Crucially, I assume that this constraint guides the interpretation of *ought*-claims: context only supplies parameters that meet the definedness constraints.²⁶ This is a fairly standard assumption about how definedness guides interpretation and one I will draw on to account for the cases in Section 2.

Ultimately my argument for this constraint is that it works: it combines with very natural assumptions about deontic and epistemic orderings to yield Agglomeration in the deontic, but not the epistemic case. Nonetheless, one might have two worries about the constraint.²⁷ First, it might appear to have substantial theoretical commitments. Do we really want to rule out from the get-go the idea that deontic *oughts* might be inconsistent? As I discuss further in Section 11, various authors have argued that deontic *oughts* can be inconsistent and so deontic dilemmas do occur.²⁸

My focus so far has been on Consistent Agglomeration, which applies only to the cases with consistent prejacent and not dilemmas. I show in Section 11 that it is possible to weaken the definedness constraint to permit dilemmas. The official stance of the paper, then, is simply that *consistent oughts* agglomerate. The case of inconsistent *oughts* is a further choice point; various options are compatible with the main ideas of the paper.

A second worry is that the definedness constraint is just tantamount to building Agglomeration into the semantics. But once we distinguish between the consistent and inconsistent version of Agglomeration, we can see this

²⁵ See Mandelkern 2019 for a different example of adding definedness conditions to modals.

²⁶ For instance, consider what the referential account of tense says about examples like:

- (i) I am happier than I was then.

Two times are salient here: the present and some past time. But there is only one reading of the sentence detectable and the reason is clear: undefinedness results if the present tense refers to the past time and the past tense refers to the present.

²⁷ Thanks to an anonymous reviewer for pushing me to consider both of these worries.

²⁸ In addition, it’s worth noting that inconsistency with dilemmas is a feature of many views of deontic modals, including the classic account.

Putting *oughts* together

is not the case. The consistency constraint does *not* entail Consistent Agglomeration — otherwise, the view could not explain the data in The Office. It is only when the constraint is combined with particular properties of the ordering, independently motivated below, that we get Agglomeration in the deontic case.

4.4 Summing up

Let's put my entries all together:

- (70) $PBEST(w, f, g, Q) = \{p \subseteq f(w) : p \neq \emptyset$
and $\neg \exists q \subseteq f(w) : q \prec_{w,f,g,Q} p\}$
- (71) a. $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q}$ is defined only if for all p and q
in $PBEST(w, f, g, Q)$, $p \cap q$ is consistent with $f(w)$.
b. If defined, $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q}$ iff $\exists p \in PBEST(w, f, g, Q) :$
 $\forall w' \in p : \llbracket \phi \rrbracket^{w',f,g,Q} = 1$

5 Relevance questions

In this section, I motivate the idea of a *relevance question*, which I propose is the contextually supplied question for my semantics; and then I discuss what particular questions get supplied.

5.1 The very idea

Certain distinctions matter in conversation and others do not. For instance, suppose we are discussing a party we attended last night and compare the following utterances:

- (72) Alice was at the party.
(73) Alice was at the party for two hours.
(74) Alice was at the party wearing size 9 shoes.

In some sense, (73) and (74) are different *ways* that (72) can be true. But they are not equally relevant: in many situations, (73) will be relevant information; only rarely will (74).

So not all ways a proposition could be true are relevant. Clearly this is driven by our interests. There are things that in most contexts we don't care

about, like someone's shoe size, their precise trajectory through space or how fast their heart is beating. For only in special cases does this information impinge on things we care about. If a proposition about Alice's appearance at the party makes distinctions amongst these kinds of things, we will tend to judge it as involving irrelevant information.

I model this notion of relevance by appeal to what I call a *relevance question*. By choosing a question whose complete answers do not distinguish between certain worlds, we can represent the relevant propositions: a proposition is relevant if it is a partial answer to the question. Otherwise it crosscuts or makes finer distinctions than the distinctions we are making in context.²⁹

5.2 The question for *ought*

I say that *ought* quantifies over answers to the relevance question. But which *particular* relevance question is given by context, when we are considering *ought*-claims? In this subsection, I'll state what I think those questions are. Later in Section 9, I'll try to put those claims on firmer ground.

In the epistemic case I take the question to be at least as fine grained as the main question which the speakers are aiming to answer. And I take this to be at least as fine-grained as the polar question corresponding to the prejacent of the *ought* claim. When considering

(75) Sarah ought to be home any minute.

²⁹ What the relationship is between a relevance question and the *question under discussion* from Roberts (2012)? If the relevance question turns out to be analysable in terms of the QUD, so much the better. But I am inclined to keep them distinct. For compare the following discourses:

- (i) A. Did Billy come to the party?
B. I'm not sure. But Alice came to the party.
- (ii) A. Did Billy come to the party?
B. I'm not sure. But Alice came to the party with a temperature of 98F.

Neither proposition is an answer to the question under discussion. But there is still a felt difference between the two: the second introduces distinctions that are in some sense not relevant at all in the context; the first does not.

In terms of possible precedents, my notion is closer to the way Stalnaker talks about possible worlds in various places. (See Stalnaker 1981, Stalnaker 1986 and Stalnaker 2014.) Hoek 2018 also employs a similar notion to make precise the notion of conversational exculpation; and Boylan & Schultheis 2021 use a similar notion in the semantics for counterfactuals.

Putting *oughts* together

the relevance question will be at least as fine-grained as the question

$\{Sarah \text{ will be home any minute}, \neg(Sarah \text{ will be home any minute})\}$

The deontic case is somewhat more complex. Here I think there are *two* kinds of relevance questions. One makes at least as many distinctions as there are available actions that we are interested in.³⁰ Take for instance the question *what did/will you do?* This question partitions the modal base into propositions that state what available action the agent performs. Take, for example, Dessert. Here the question *what will you do?* is the partition below:

just pie	just cannoli
just cheesecake	pie + cannoli
pie + cheesecake	neither
cannoli + cheesecake	all three

Each total outcome is represented by its own proposition.

The other kind of question simply asks *How good will the action you perform be?* and does not distinguish between equally good actions. This question partitions the modal base into propositions which say that the agent performed one of the best actions, one of the second best actions, and so on. In general, the partition will look something like this:

$\{I \text{ do a best action}, I \text{ do a second best action}, \dots, I \text{ do a worst action}\}$

Given the information in Dessert, the partition is contextually equivalent to:

³⁰ For evaluative rather than deliberative uses of *ought*, these might instead be outcomes, rather than available actions.

just pie or just cannoli	
just cheesecake	pie + cannoli
pie + cheesecake	or neither
or cannoli + cheesecake	all three

The different kinds of question reflect different ways we can use actions to distinguish between possibilities. *What will you do?* distinguishes between them exactly as finely as our options do. *How good will what you do be?* only distinguishes between possibilities up to how good the our actions are, lumping possibilities with equally good actions into one cell. As a result, *What will you do?* tends to partition more finely than *How good will the action you perform be?*

Why think that these questions can be the relevance question? Because the relevance question is informed by our interests; and, when thinking about decision-making, these are the kinds of questions we care to know the answer to. We want to know what choices people will make; and we want to know how good their choices will be. As we will see, this predicts Indifference by filtering out of an ordering the propositions that violate the consistency constraint.

6 Questions and orderings

The apparatus of questions allows us to state the fundamental difference between epistemic and deontic orderings, the one crucial for securing deontic and not epistemic Agglomeration.

Suppose that p is the disjunction of some elements q_1, \dots, q_n of a question Q ; the q 's are the relevant *ways* in which p . Must some way for p to be true be at least as good as p itself? Or could the q 's all be *worse* than p ?

Putting *oughts* together

For epistemic value, the answer is yes: a proposition can be epistemically better than its constituent cells. For example, if epistemics are based on probability, then p can clearly be ranked higher than any of its constituent Q -cells: a disjunction will usually be more probable than either of its disjuncts. The same applies to expectation and normality: I might expect a fair coin to land on either heads or tails, without expecting it to land on a particular side; likewise *heads or tails* obtains in a greater proportion of normal situations than simply *heads*.

With deontics I submit that it is different: p cannot be better than *all* of its constituent Q -cells. Think of q_1, \dots, q_n as the ways that p can come about. p may well be better overall than many or even most of its realisations. But p cannot be better than *all* of its possible realisations: at least one Q -cell must be at least as good as p . Likewise, it cannot be *worse* than all ways for it to come about: its value must lie somewhere between the best and worst ways for it to come about.

Why think this? First, deontic orderings that violate this constraint seem quite defective. Consider some sentences that contradict the constraint:

- (76) I prefer pie to ice cream. But I prefer ice cream to cherry pie, apple pie, blueberry pie... In fact, I prefer ice cream to any particular kind of pie.
- (77) It's morally better to give \$1,000 to a malaria charity than to UNICEF. But giving it to UNICEF is better than giving it to any particular malaria charity.

Both of these sentences say that some A is better than some B , even while B is better than any particular way C for A to obtain: in the first example A is having pie and B is having ice cream and the C s are having apple pie, blueberry pie and so on; in the second, A is giving the \$1k to a mosquito net charity, B is giving it to UNICEF and the C s are the particular malaria charities. So both examples violate the constraint and are much the worse for it.

Second, this constraint is entailed by the popular claim that the deontic value of a proposition should be a *weighted average* of the values of the different ways for it to be true. First take simple expected utility theory in the style of [Savage 1972](#), where the value of an action is the sum of the values of the outcomes of that action weighted by their probabilities. Say that S is the set of states the world might be in, $O(A, s)$ is the outcome A produces

when s obtains and $V(o)$ is the value of an outcome. We can then write the expected value of A as follows:

$$EV(A) = \sum_{s \in S} Pr(s)V(O(A, s))$$

Here it is a straightforward mathematical fact that the expected value of A must fall somewhere between the values of best and worst possible outcomes it might lead to.

Basically all serious views in decision theory follow suit here: the value of a proposition as a weighted average of the ways it could be true. Other expected value decision theories agree with Savage's theory here: both evidential and causal decision theory entail that preferences should have the ordering property that I state above.³¹ Even the relative newcomer of risk-weighted expected utility, formulated by Buchak (2013) and there argued to better capture certain patterns of risk aversion, also delivers this property. Since there is such widespread agreement on the deontic constraint, I think it is plausibly deontically neutral, in the sense defined in Section 4.

Summing up, we have:

If g is deontic and q is a partial answer to Q , then, where p , $p' \subseteq q$ are respectively best and worst complete answers that entail q , $p \succsim_{w,f,g,Q} q \succsim_{w,f,g,Q} p'$.

If g is epistemic and q is a partial answer to Q , then there may be no complete answer $p \subseteq q$ such that $p \succsim_{w,f,g,Q} q$.

This difference in deontic and epistemic orderings is exactly why Agglomeration fails only in the epistemic case.

7 Deontic *ought* is a box after all

Before spelling out the predictions, it will be helpful to note my view has the following feature:

If defined, 'ought ϕ ' is true iff it is entailed by the best complete answer to the relevance question.

³¹ For evidential decision theory, see Jeffrey 1965, ch.9: the deontic constraint follows from axioms 3(a) and (b). For causal decision theory, see Joyce 1999 ch.7; since causal decision theory also has Jeffrey's axioms 3(a) and (b), the same argument applies.

This means that, when defined, the deontic reading of *ought* essentially has the truth-conditions of the classic semantics.³²

Why is this? First, we can show there is always a unique complete answer which is deontically best, when our selection of parameters yields definedness. Suppose p is a best partial answer.³³ Since it is a partial answer, p must be entailed by some complete answer q . The deontic constraint ensures *at least one* of the complete answers entailing p is also best: otherwise p would be better than *any* of the ways for it to be true. But also *at most one* complete answer can be best, if ‘ought ϕ ’ is defined: complete answers are disjoint and so having multiple best complete answers would violate the consistency constraint. So exactly one complete answer is best.

But this means that, when any deontic *ought*-claims are defined, what deontically ought to be the case is simply whatever the best answer to the relevance question entails. Any best partial answer not entailed by the best complete answer would be inconsistent with it. So any best proposition must be one that is entailed *already* by the best complete answer. It is a direct consequence of this that deontic Agglomeration must hold: if ‘ought ϕ ’ and ‘ought ψ ’ are true on a deontic reading, then some best complete answer must entail both ϕ and ψ .

Importantly, no such thing holds for epistemics. Generally *none* of the complete answers will be *epistemically* best. This is perfectly coherent, given the epistemic constraint: merely partial answers will tend to be epistemically better than the complete answers that entail them. The epistemic *ought* will *not* have the truth-conditions of the classic semantics, whenever defined; and for this reason epistemic Agglomeration fails.

8 Predictions

Now we are in a position to explain our data. First of all, we will see that epistemic Agglomeration fails: the semantics permits Agglomeration failure in principle; and the consistency constraint permits epistemic orderings that violate Agglomeration, given the generalisation about epistemic orderings in Section 6. We will then see that deontic Agglomeration does not fail: deontic orderings which violate Agglomeration are *not* permitted by the consistency constraint, given the generalisation about deontics in Section 6. Finally, we will see that Indifference is predicted to hold: in cases with multiple best

³² The facts below are proved in the Appendix.

³³ Note that the Limit Assumption entails there must be at least one.

options, the consistency constraint rules out the relevance question *what will I do?*, while the other salient relevance question *How good an option will I perform?* forces a disjunctive *ought* to hold.

To get a sense of how these predictions are obtained, I will first work through the particular cases we have seen so far; then I will zoom out and state some more general predictions.

8.1 Predictions for epistemics

Let's start by thinking about what relevance question Q and ordering $\lesssim_{w,f,g,Q}$ are plausible in The Office.

The relevance question will be *which workers are in?*: ultimately, what we are interested in here is what particular workers are in. The relevance question will thus have as its complete answers propositions like *Alice is in the office, Billy is in the office, Carol is in the office, ..., and Zadie is in the office*; or *Alice is not in the office, Billy is in the office, Carol is not in the office, ..., and Zadie is in the office*; or *Alice is in the office, Billy is not in the office, Carol is not in the office, ..., and Zadie is not in the office*. There will be a complete answer for each such possible combination of workers that might be in today.

I assume the ordering of partial answers to the relevance question is a function of their levels of expectation or normality; and we said that in The Office specifically, these relations appear to coincide with degrees of probability. I assume then that in The Office the set of best propositions will be those that pass some threshold probability greater than 0.5 probability, so that we have the following ordering.

$p \lesssim_{w,f,g,Q} q$ iff one of the following conditions holds:

- i. $P(p \mid \bigcap f(w)) > 0.5$; or
- ii. $P(p \mid \bigcap f(w)) \geq P(q \mid f(w))$

This ordering forces those above the 0.5 threshold to be best: if p is above the threshold, then automatically for any other q , $p \lesssim_{w,f,g,Q} q$ and so p is best. This is as it should be: p may epistemically ought to be the case, even if it does not have maximal probability. Below the threshold, the ordering tracks probability directly: $p \lesssim_{w,f,g,Q} q$ iff p is at least as probable as q .³⁴

³⁴ As an anonymous reviewer correctly points out, we do not want our judgements in this case to be overly sensitive to small fluctuations in probability. This will hold on such an ordering, provided the relevant propositions are not too close to the threshold.

Putting *oughts* together

Such a threshold will mean the set of *best* propositions must include the following set of partial answers:

{Alice is in the office today, Billy is in the office today, ... , Zadio is in the office today, Not everybody is in the office today}

On the other hand, none of the complete answers will be best, as their probability/normality will be too low. In fact, $PBEST(w, f, g, Q)$ is guaranteed to be pairwise contextually-consistent: it is a simple consequence of probability theory that p and q cannot *both* have greater than 0.5 probability, if they are inconsistent.

This selection of parameters correctly predicts the truth of

- (2) Alice should be in the office today.
- (3) Bob should be in the office today.
- (4) Carol should be in the office today.
- ...
- (5) Zadio should be in the office today.

and the falsity of

- (6) Everyone should be in the office today.

There is a best proposition that entails Alice will be in the office, that entails Bob will be the office, and so on. But no best proposition entails everybody will be in the office.³⁵

Notice too that this particular selection of parameters does not generate any dilemmas: we do not *also* predict the truth of

- (66) Alice should *not* be in the office today.

³⁵ We can also see why many conjuncts are required for Agglomeration to fail. Take for instance:

- (i) Alice and Bob should be in the office today.

The claim that Alice and Bob are both in the office today will have high probability. Indeed, in this example, it is quite plausible the two conjuncts are close to probabilistically independent, giving the conjunction probability of around 0.93. Simple probability calculations show that the number of conjuncts will have to be relatively large before the probability drops below 0.5.

No best proposition can entail this, given that the threshold is above 0.5. Moreover, *no* selection of parameters can predict the simultaneous truth of (6) and (66), as doing so would violate the consistency constraint.

8.2 Predictions for deontics

Now recall Dessert:

Dessert. There are three dessert options: cannoli, cheesecake, and apple pie. Pie and cannoli are tastiest. I can order as many dishes as I like, but I will definitely feel ill if I have more than one.

I argued that in such cases there are two possible relevance questions, *what will I do?* and *how good will the action I perform be?* Only the latter is a possible question parameter in this context; and it yields Indifference.

Take *what will I do?* first. We saw this gives us the partition:

just pie	just cannoli
just cheesecake	pie + cannoli
pie + cheesecake	neither
cannoli + cheesecake	all three

The natural value for g will simply be a function which tracks my preferences: $\succsim_{w,f,g,Q}$ ranks the partial answers to Q based on my actual preferences, given the background information.

Given the set-up of the case, the ordering of partial answers based on this partition will be inconsistent. *I have just pie* and *I have just cannoli* will be among the best propositions, given the question above. But we saw in

Putting *oughts* together

Section 7 that, whenever deontic *ought*-claims are defined, there must be a unique best complete answer to the relevance question. This question yields undefinedness and thus no possible interpretation will result from this relevance question.

Assuming the ordering tracks my preferences, we can further argue that there simply can be *no* reading of (17) and (18) where they are both true.

(17) I ought to have pie.

(18) I ought to have cannoli.

On any admissible set of parameters where (17) and (18) were *both* true, some uniquely best complete answer would have to entail that I have pie *and* cannoli. But such a proposition could never be best: I prefer to have one dessert rather than two.

Let's now take the other relevance question, *how good will the action I perform be?*, and see why it predicts Indifference holds good here. We saw this question induces the following partition on the modal base:

just pie or just cannoli	
just cheesecake	pie + cannoli
	or
pie + cheesecake	neither
or	
cannoli + cheesecake	all three

If \succsim is my preference ordering, the proposition *I have pie or cannoli* will be among the best propositions. This gives us the truth of

(78) I ought to have pie or cannoli.

This is exactly what we wanted.³⁶

8.3 Zooming out

To sum up, I think it's helpful to zoom out and say how the pieces of my view work together.

The different properties of deontic and epistemic orderings give us the different verdicts on Agglomeration. In the deontic case, when *ought*-claims are defined, there must be a unique complete answer which is among the best propositions; and whether *ought*-claims are true depend simply on what that answer entails. In the epistemic case, there may be no complete answer among the best propositions. The deontic and epistemic constraints interact with the definedness condition to ensure this. In the Appendix I prove:

Fact 4. (No deontic Agglomeration Failure) There are no w, f, g, Q such that:

- i. $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q} = 1$ and $\llbracket \text{ought } \psi \rrbracket^{w,f,g,Q} = 1$
- ii. g , the ordering function, is deontic;
- iii. and $\llbracket \text{ought } (\phi \wedge \psi) \rrbracket^{w,f,g,Q} = 0$

Fact 5. (Epistemic Agglomeration Failure) There are w, f, g and Q such that

- i. $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q} = 1$ and $\llbracket \text{ought } \psi \rrbracket^{w,f,g,Q} = 1$
- ii. g , the ordering function, is epistemic;
- iii. and $\llbracket \text{ought } (\phi \wedge \psi) \rrbracket^{w,f,g,Q} = 0$

³⁶ As the editor points out, the notion of an ordering is invoked twice here. To avoid any appearance of circularity, here is how to state the construction more carefully. We are defining the outputs of g , a function from worlds, questions and modal bases to an ordering over answers to the supplied question. When supplied with the question *What will I do?* — call this question Q —, I assume g outputs the natural ordering over the options. I use this output to construct the cells of a different question *How good will the action I perform be?*: this question — call it Q' — is

$$\{p \subseteq f(w) : \text{for some } S \subseteq Q, p = \bigcup S \text{ and for any } p', p'' \in S, p' \approx_{w,f,g,Q} p''\}.$$

Q' is fed back into g , along with w and f , to construct a new ordering, $\lesssim_{w,f,g,Q'}$.

The relevance question enters the analysis in two places. First, it allows us to state the epistemic and deontic constraints. Second, it secures Indifference. We know that, given a question with multiple best answers, my semantics yields undefinedness. But this does not yet predict that disjunctive *oughts* are true in cases like Dessert. It is because *multiple* questions are available, in particular the question *how good will the action I perform be?*, that I predict Indifference.

9 Constraining the Relevance Question

We have seen how my view predicts the facts, given the assumptions from Section 5 about the relevance question. In this section, to put those assumptions on firmer ground, I will motivate a principle governing how the relevance question evolves.³⁷

Relevance is of course a context-dependent matter. And it's plausible that the prejacent of a modal can change the relevance question. Suppose I say:

(79) Alice must have been at the party for three hours.

This will tend to make the length of time relevant: after an utterance of (79), the relevance question will tend to refine the question *how long did Alice spend at the party?*³⁸ More generally, when there is a set of salient alternatives to the prejacent of the modal, distinguishing between these alternatives will tend to become relevant.³⁹

Still, the relevance question will sometimes be harder to update, because of the definedness conditions. Recall Dessert, where we have a number of equally good alternatives. If the relevance question is simply refined with the alternatives to having pie, for instance, the resulting relevance question will lead to undefinedness. This rules out the most natural generalisation, namely that we simply refine the prior relevance questions with the modal's prejacent and its alternatives.

³⁷ Thanks to the editor for pushing me to say more here.

³⁸ Note that existing distinctions may continue to be relevant, which is why the relevance question should refine the question introduced by the modal, rather than simply be replaced by it. (Recall that Q refines Q' iff when $q \in Q$, there is some q' in Q' such that $q \subseteq q'$.)

³⁹ These alternatives might be supplied by the focus structure of the prejacent of the modal; see Rooth 1992, Schwarzschild 1999 and Beaver & Clark 2008. I assume that in deontic cases, the alternatives are generally the agent's options.

I propose accommodation seeks a compromise between these competing pressures. Uttering a modal will tend to make a set of alternatives relevant. But accommodation should also seek to make our utterances defined. A natural idea then is that accommodation looks for a way to refine the relevance question that tracks the salient alternatives as closely as it can, while still yielding definedness. The following rule makes this idea more precise:

Accommodation Rule. Suppose that c is the prior context and that $\llbracket \text{ought } \phi \rrbracket^{w_c, f_c, g_c, Q_c}$ is defined. Then c' , the context after uttering $\lceil \text{ought } \phi \rceil$ has as its relevance question some $Q_{c'}$ such that:

- i. $\llbracket \text{ought } \phi \rrbracket^{w_{c'}, f_{c'}, g_{c'}, Q_{c'}}$ is defined;
- ii. $Q_{c'}$ is refined by the result of refining the prior relevance question with the alternatives to ϕ ;
- iii. there's no other Q' that properly refines $Q_{c'}$ while also meeting i and ii.


In many cases, this rule tells us exactly what the new relevance question should be. When refining Q_c with the alternatives to ϕ yields definedness, the resulting question is the new relevance question. Since epistemics do not violate the consistency constraint, the relevance question in general is a refinement of the polar question corresponding to the prejacent. That is, the new relevance question is just the old one refined with the alternatives introduced by the modal.

In other cases, the rule does not pin down a unique relevance question, but still greatly constrains the possible values. In the deontic case there is no guarantee of definedness, when we refine with the alternatives supplied by ϕ . Thus the accommodation rule does not always pin down a unique value here. But it does ensure accommodation will search for a new relevance question that disjoins the best complete answers to the prior relevance question refined with ϕ and its alternatives.

I show in the Appendix that given the stronger version of my consistency constraint from footnote 24, the question *how good an action will I perform?* must be a possible value for the relevance question. That question results from disjoining equally good options, satisfying consistency; but it also does not disjoin more complete answers than it has to, meaning no other question satisfying the Accommodation Rule refines it.

Now, this question does not *uniquely* satisfy the Accommodation Rule. For example, in Dessert the following partition would also satisfy the rule:

Putting *oughts* together



But I assume the Accommodation Rule is not the only rule governing how we accommodate. Another plausible constraint is that the accommodated partition correspond to some natural question we could actually ask. The partition above does not; nor, as far as I can see, does any other way of satisfying the rule. Thus, I think that this Accommodation Rule bears out the claims of Section 5.2.

10 Patching up the competition?

An anonymous reviewer asks whether other views can use my machinery to achieve similar results.

For the classic semantics and the conflict account, the answer is no. Agglomeration is simply hard-wired into the classic semantics. Though the conflict account does not validate Agglomeration, I think a similar problem holds. As I showed in Section 4.3, the conflict account incorrectly entails *any* Agglomeration failure involves a dilemma.

Contrastivism has the best shot at mimicking my strategy. To do so, the contrastivist could identify their alternatives with elements of the relevance question. Define $ALT(Q, p)$ to be the $\neg p$ -entailing elements of Q ; then we could restate contrastivism as follows:

- (80) a. $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q}$ is defined only if $PBEST(w, f, g, Q)$ is pairwise consistent;
- b. if defined, $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q} = 1$ iff for all $q \in ALT(\phi)$: $\llbracket \phi \rrbracket^{f,g,Q} > q$

This view appears to validate deontic Agglomeration but not epistemic Agglomeration, if we supplement it with my pragmatics.

Given just the aims of this paper, I don't see that there is much to choose between my view and this revised contrastivism. My main innovation is leveraging the definedness constraint and the ordering properties together to yield different logics for deontics and epistemics. The existential semantics is a good launching pad for these pieces; but perhaps not the only one.

In the broader dialectic, though, I think the existential semantics has an important advantage. Contrastivist semantics generally aim to invalidate the principle of Inheritance:

Inheritance: If $\phi \models \psi$, then $\lceil \text{ought } \phi \rceil \models \lceil \text{ought } \psi \rceil$

Cariani (2016a) shows that, for this reason, many contrastivist semantics invalidate a weaker, but very plausible, principle he calls Weakening:

Weakening: $\lceil \text{ought } \phi \rceil, \lceil \text{ought } \psi \rceil \models \lceil \text{ought } (\phi \vee \psi) \rceil$

The modified contrastivist semantics above also invalidates Weakening: Cariani's counterexamples arise for agents whose preferences align with expected utilities. I agree with Cariani that this is a serious problem.

Cariani (2013b,a, 2016b) develops the only contrastivist semantics I know of that invalidates Inheritance while validating Weakness; but that semantics validates Agglomeration in full generality. I leave it to the contrastivist to develop a view that validates deontic Agglomeration and Weakening but not Inheritance or epistemic Agglomeration.

11 Dilemmas

Consistent Agglomeration has been our focus. What of Inconsistent Agglomeration?

Inconsistent Agglomeration: if ϕ and ψ are *inconsistent*, then $\lceil \text{ought } \phi \rceil, \lceil \text{ought } \psi \rceil \models \lceil \text{ought } (\phi \wedge \psi) \rceil$

My theory also validates this too. But some think moral dilemmas are counterexamples to this inference. Take the following case:⁴⁰

⁴⁰ Sinnott-Armstrong (1985) calls this a *symmetric dilemma*.

Putting *oughts* together

Sophie's Choice. Sophie is forced to choose which of her children is going to be sent to the labour camp and which is going to be killed. If she chooses neither child, then both will be killed.

Here there is a true reading of both of the following.

(81) Sophie ought to save her daughter.

(82) Sophie ought to save her son.

But clearly it is not true that:

(83) Sophie ought to save her daughter and her son.

For Sophie cannot do both of these things together. So, if all of these are evaluated in the same context, we have a failure of Agglomeration.

The topic of dilemmas has its own rich literature. Rather than give the last word, I aim to show my theory has good options to choose from. One is to relax the consistency constraint just enough for Inconsistent Agglomeration to fail. Another is to argue on independent grounds that dilemmas are not genuine counterexamples to Inconsistent Agglomeration. Both I think are defensible options.

11.1 Modifying the Consistency Constraint

One option is to weaken the consistency constraint by appeal to incomparability.

As von Fintel (2012) notes, dilemmas do not simply arise whenever we have more than one best option; cases like Dessert are clearly not dilemmas. Following van Fraassen (1973), Horty (2012) and Swanson (2016), I suggest that in dilemmas the best options are *incomparable*.⁴¹

Why think this? Because *mild sweetening* of an option does not resolve the dilemma.⁴² Suppose Sophie can choose between having her daughter saved and having her son saved *and* treated slightly better in the labour camp. Having her son saved *and* treated slightly better in the labour camp is better than just having her son saved; but improving this option does not resolve the dilemma. So saving her son and saving her daughter cannot be equally

⁴¹ I mean this simply in the sense that neither is at least as good as the other.

⁴² The name is from Hare (2010).

good. Neither option is *better* than the other, so we should conclude the two options are incomparable.

In light of this, we could weaken the consistency constraint to say merely each pair of *comparable* best propositions must be consistent:

Comparable Consistency: $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q}$ is defined only if for every *comparable* p and q in $PBEST(w,f,g)$, $p \cap q \cap \cap f(w) \neq \emptyset$.

In Sophie’s Choice, the best propositions are *Sophie saves her daughter* and *Sophie saves her son*. While contextually inconsistent, they are not comparable. Comparable Consistency permits this ordering to produce a possible true interpretation of (81) and (82).⁴³

This is one way I can accommodate dilemmas, but not obviously the only one. Goble (2005) notes that we might accommodate dilemmas by weakening the rule of Inheritance. Finally, von Fintel (2012), Gillies (2012) and Horty (2014) all note that the classic semantics can be generalised in various ways to accommodate deontic conflict. Given that, on my theory, deontic *oughts* have the truth-conditions of the classic semantics, there might be similar ways of generalising my semantics — for instance, by quantifying over different orderings — to accommodate dilemmas.⁴⁴

11.2 Rejecting dilemmas

Alternatively context-sensitivity might explain the data in Sophie’s Choice. I give a new argument here for this approach, one based on the *pseudo-factivity* of deontic *must*.

As observed by von Fintel (2012), apparent dilemmas arise for *must* and *have to* as well as *ought*. In Sophie’s Choice both of the following seem to have true readings:

(84) Sophie must/has to save her daughter.

⁴³ Note this validates Consistent Agglomeration only for comparable ϕ and ψ : Goble (1996) shows even Consistent Agglomeration is inconsistent with the existence of dilemmas. I think this is a plausible restriction for those that think dilemmas are genuine, given its similarity to the version of Agglomeration validated by the systems in Horty 2012 and van Fraassen 1973. McNamara (2004) suggests an interesting alternative way to restrict Agglomeration.

⁴⁴ Goble (2013) notes that another approach is to weaken the rule of Substitution of Logical Equivalents in *oughts*. I am less optimistic that this move could be recapitulated in my semantics.

Putting *oughts* together

(85) Sophie must/has to save her son.

This is important: if there are dilemmas, they can be expressed using strong deontic necessity too. So any account of dilemmas, or the appearance of dilemmas, must apply to *must* and *have to*.

But there is good reason to think that these apparent dilemmas for *must* are not genuine. As Sinnott-Armstrong (1985) and Horty (2012) note, dilemmas yield contradictions assuming natural principles for deontic *must* and *may*:

Duality: $\models \text{`must } \phi \leftrightarrow \neg \text{may } \neg \phi \text{'}$

Must-to-may: $\text{`must } \phi \text{' } \models \text{`may } \phi \text{'}$.

Must-Inheritance: If $\phi \models \psi$, $\text{`must } \phi \text{' } \models \text{`must } \psi \text{'}$

In a genuine dilemma, for some inconsistent ϕ and ψ , $\text{`must } \phi \wedge \text{must } \psi \text{'}$ is true. Given the above, this entails $\text{`may } \neg \phi \wedge \neg(\text{may } \neg \phi) \text{'}$, which of course cannot be. Some assumption must be rejected; the existence of dilemmas looks to me to be the weakest.

Context-sensitivity is the obvious tool to reach for. If (84) is evaluated in a different context, and so against a different ordering, from (85), then we avoid contradiction without denying any of the principles above.⁴⁵ Naturally, this explanation would extend to *ought* too, relieving the pressure on my consistency constraint.

There is also an independent reason to favour context-shifting. *Must* is *pseudo-factive*: as Ninan (2005) observes, it sounds incoherent to say

(86) #You must clean your room, even though you aren't going to.

Presumably deontic *must* does not *entail* its prejacent; but it still commits the speaker to the prejacent. This makes the existence of dilemmas for *must* surprising. If *both* (84) and (85) are true, then, by pseudo-factivity, there should be a felt commitment to both Sophie saving her daughter and saving her son. Since both (84) and (85) are acceptable, there can be no such felt commitment.

I suggest this is further evidence for context-sensitivity. Notice when it is clear that multiple deontic orderings are relevant in the context, the pseudo-factivity of *must* disappears. If I say something like:

⁴⁵ Something like this approach is suggested by Castaneda 1981. Brink (1994) is also sympathetic to thinking that genuine dilemmas do not exist.

(87) According to your father, you must be home by 9; and according to your mother you must be in bed by 10.

Neither of my *must* claims here have the usual felt entailment. I can easily continue with:

(88) But of course, you're not going to do either of those things.

When there are multiple sources of obligation in play and we have not committed ourselves to any one of them, *must* is not pseudo-factive.

Now if (84) and (85) were interpreted relative to the same ordering we should feel committed to contradictory propositions. But if there were multiple orderings in play, neither completely endorsed or taken to be binding, we predict no such commitments. Since the latter is what we actually see, the multiple ordering strategy does well here.

An anonymous reviewer notes that (87) is reportative, rather than directive: it states that a *must*-claim holds, but does not itself issue a directive. They suggest this might explain the absence of pseudo-factivity. But, as [Ninan \(2005\)](#) observes, even reportative uses of *must* are pseudo-factive. Suppose that Alice says of her older brother Billy:

(89) #He must do his homework; but he is not going to.

The *must* here is not directive: Alice is not even addressing Billy. But still it is pseudo-factive.

12 Conclusion

To explain why epistemic but not deontic Agglomeration fails, I gave a new semantics where *ought* is an existential quantifier over best propositions and so $\lceil \text{ought } \phi \rceil$ and $\lceil \text{ought } \psi \rceil$ can be true, while $\lceil \text{ought } (\phi \wedge \psi) \rceil$ is false. I added a layer of question-sensitivity to *ought*, so that the best propositions must also be partial answers to a background question; I also added a pairwise consistency constraint. Together with some assumptions about the pragmatics of the background question and orderings, these deliver deontic but not epistemic Agglomeration.

Putting *oughts* together

Appendix

We make the following assumptions about orderings:

Assumption 1 (Limit assumption). For every w, f, g, Q , for any non-empty set of propositions S there is some $p \in S$ such that $p \neq \emptyset$ and for no $p' \in S$ $p' \prec_{w,f,g,Q} p$.

Assumption 2 (Deontic orderings). If g is deontic, then if p is a partial answer to $Q|f(w)$, then, where $q, q' \subseteq p$ are respectively best and worst complete answers that entail p , $q \succsim_{w,f,g,Q} p \succsim_{w,f,g,Q} q'$.

Assumption 3 (Epistemic orderings). If g is epistemic and p is a partial answer to $Q|f(w)$, then there may be no complete answer $q \subseteq p$ such that $q \succsim_{w,f,g,Q} p$.

Assumption 4 (Question-sensitivity). All g 's are *question-sensitive* ordering functions, i.e. $\prec_{w,f,g,Q}$ orders only partial answers to Q .

Assumption 5 (Partition invariance). If $Q \leq Q'$ and q, q' are partial answers to Q' , $q \succsim_{w,f,g,Q} q'$ iff $q \succsim_{w,f,g,Q'} q'$.

Assumption 6 (Comparability). If p is a partial answer to $Q|f(w)$, then for all $q \in Q|f(w)$ such that $q \subseteq p$ and q are comparable, i.e. either $p \succsim_{w,f,g,Q} q$ or $q \succsim_{w,f,g,Q} p$.⁴⁶

The following fact is helpful and easily proved:

Fact 1. If p is a partial answer to $Q|r$ and q is a complete answer, then either $q \subseteq p$ or $p \cap q = \emptyset$.

We now prove:

Fact 2. If g is deontic and obeys the consistency constraint with respect to w, f and Q , then there is exactly one q such that $q \in Q|f(w)$ and $q \in PBEST(w, f, g, Q)$.⁴⁷

⁴⁶ $Q|r$ is the restriction of Q to r i.e. $\{s: \exists q \in Q : s = q \cap r\}$

⁴⁷ Note that g obeys the consistency constraint with respect to w, f , and Q iff for some ϕ $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q} \neq \#$.

Proof. Suppose g is deontic. By the Limit Assumption, there is some $p \in PBEST(w, f, g, Q)$. By question-sensitivity, p must be a partial answer to $Q|f(w)$. So there must be some $q \in Q|f(w) : q \subseteq p$. Assume for contradiction, that for all $q \in Q|f(w)$: if $q \subseteq p$ then $q \notin PBEST(w, f, g, Q)$. By the Limit Assumption there must be some best q' such that $q' \subseteq p$. It then follows that p is strictly better than q' . But this violates the deontic constraint. So for some p -entailing $q \in Q|f(w)$: $q \in PBEST(w, f, g, Q)$. Uniqueness is secured by the consistency constraint: since complete answers are inconsistent, the consistency constraint would be violated, if there were more than one such q' .

Fact 3. (Deontic ought is boxy) If $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q} \neq \#$ and g is deontic, then there's some $q \in Q$ such that $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q} = 1$ iff $q \subseteq \llbracket \phi \rrbracket^{f,g,Q}$.

Proof. Suppose the antecedent holds. By Fact 2, we know that there is a unique best $q \in Q|f(w)$; we show that $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q} = 1$ iff $q \subseteq \llbracket \phi \rrbracket^{f,g,Q}$.

The right-to-left direction is obvious, since $q \in PBEST(w, f, g, Q)$. So assume $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q} = 1$. Then there must be some p that is a partial answer to Q , is an element of $PBEST(w, f, g, Q)$ and which entails $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q} = 1$. Since p is a partial answer to Q and q is a complete answer, we know by Fact 1 that either $q \subseteq p$ or $q \cap p = \emptyset$. But if the latter obtained, the consistency constraint would be violated. So q entails p .

Fact 4. (No deontic Agglomeration Failure) There are no w, f, g, Q such that:

- i. $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q} = 1$ and $\llbracket \text{ought } \psi \rrbracket^{w,f,g,Q} = 1$
- ii. g is deontic;
- iii. and $\llbracket \text{ought } (\phi \wedge \psi) \rrbracket^{w,f,g,Q} = 0$

Proof. If (a) and (b) hold, then it is immediate that (c) does not, given Fact 3.

Fact 5. (Epistemic Agglomeration Failure) There are w, f, g and Q such that

- i. $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q} = 1$ and $\llbracket \text{ought } \psi \rrbracket^{w,f,g,Q} = 1$

Putting *oughts* together

- ii. g is epistemic;
- iii. and $\llbracket \text{ought } (\phi \wedge \psi) \rrbracket^{w,f,g,Q} = 0$

Proof. The selection of parameters given in Section 8.1 suffices.

Finally, we prove the result mentioned in Section 9. First we recursively define the n th best propositions. If S is a set of propositions, say that $\neg S = \{\neg p : p \in S\}$. Then:

$$PBEST_1(w, f, g, Q) = PBEST(w, f, g, Q)$$

If $PBEST(w, f + \bigcup(\neg PBEST_n(w, f, g, Q)), g, Q) \neq \emptyset$,

then $PBEST_{n+1}(w, f, g, Q) =$

$$PBEST(w, f + \bigcup(\neg PBEST_n(w, f, g, Q)), g, Q);$$

otherwise $PBEST_{n+1}(w, f, g, Q)$ is undefined.

Note that for epistemics only $PBEST_1$ defined, since the tautology is among the epistemically best propositions. This is good: we don't want our constraint to require the set of complete answers to be equally *epistemically* good.

Given a question, a modal base and an ordering function, we define the question that disjoins the n th best complete answers to Q . Say that:

$$s_n = \bigcup\{q \in Q : q \in PBEST_n(w, f, g, Q)\}$$

$$S_{Q,f,g} = \{q : q = s_n \text{ for some } n\}$$

$S_{Q,f,g}$, is the desired question.

Recall now the Strong Consistency Constraint from footnote 24:

Strong Consistency Constraint: $\llbracket \text{ought } \phi \rrbracket^{w,f,g,Q}$ is defined only if for all p and q in $PBEST_n(w, f, g, Q)$, $p \cap q$ is consistent with $f(w)$, for all n such that $PBEST_n(w, f, g, Q)$ is defined.

Say that $\phi_c?$ is the contextually salient partition supplied by ϕ and we can state the Accommodation Rule from Section 9 as:

Accommodation Rule. Suppose that c is the prior context and that $\llbracket \text{ought } \phi \rrbracket^{w_c, f_c, g_c, Q_c}$ is defined. Then c' , the context after uttering $\lceil \text{ought } \phi \rceil$ has as its relevance question some $Q_{c'}$ such that:

- i. $\llbracket \text{ought } \phi \rrbracket^{w_c, f_c, g_c, Q_{c'}} \neq \emptyset$;
- ii. $Q_c + ALT(\phi, c) \leq Q_{c'}$;
- iii. there's no other Q' that properly refines $Q_{c'}$ while also meeting i and ii.

Given our above definition, $S_{Q_c + \phi_c?, f_c, g_c}$ is the result of taking the prior question, refining it with $\phi_c?$ and then disjoining the n th best complete answers, for every n . We prove:

Fact 7. Given the Strong Consistency Constraint, when g is deontic, $S_{Q_c + \phi_c?, f, g}$ satisfies the Accommodation Rule.

Proof. First, note that the Strong Consistency Constraint is the only source of indeterminacy in our language so condition 1 will be met iff the consistency constraint obtains.

Take an arbitrary n where $PBEST_n(w, f, g, S_{f, g, Q_c + \phi_c?})$ is defined. First we show $PBEST_n$ contains a unique complete answer, call it s_n . Being defined, it contains some proposition; and by the deontic constraint it contains some complete answer to $S_{Q_c + \phi_c?, f_c, g_c}$. By the definition of $S_{Q_c + \phi_c?, f_c, g_c}$, any other distinct complete answer must be identical to some s_i , where $i \neq n$. So either $i < n$ or $n < i$.

In the former case, s_i is inconsistent with $f + \bigcup(\neg PBEST_{n-1}(w_c, f_c, g_c, Q_c + \phi_c?))(w_c)$ and so cannot be an element of $PBEST_n$. So suppose $n < i$. By definition, s_n and s_i are the unions of the elements of $PBEST_n(w, f, g, Q_c + \phi_c?)$ and $PBEST_i(w, f, g, Q_c + \phi_c?)$ respectively. By the Deontic Constraint and Comparability, if $q_n \in PBEST_n(w, f, g, Q_c + \phi_c?)$ then

$$s_n \approx_{w, f + \bigcup(\neg PBEST_{n-1}(w_c, f_c, g_c, Q_c + \phi_c?)), g, Q_c + \phi_c?} q_n$$

and there is some $q_n \in PBEST_n(w, f, g, Q_c + \phi_c?)$ such that for any $q_i \in PBEST_i(w, f, g, Q_c + \phi_c?)$:

$$q_n \prec_{w, f + \bigcup(\neg PBEST_{n-1}(w_c, f_c, g_c, Q_c + \phi_c?)), g, Q_c + \phi_c?} q_i$$

But then, by the deontic constraint and transitivity,

$$s_n \prec_{w, f + \bigcup(\neg PBEST_{n-1}(w_c, f_c, g_c, Q_c + \phi_c?)), g, Q_c + \phi_c?} s_i.$$

Putting *oughts* together

Then by Partition Invariance,

$$s_n \prec_{w, f + \cup(\neg PBEST_{n-1}(w_c, f_c, g_c, Q_c + \phi_c?), g, S_{f, g, Q_c + \phi_c?})} s_i$$

and so $s_i \notin PBEST_n(w, f, g, S_{f, g, Q_c + \phi_c?})$.

We now prove $PBEST_n(w, f, g, S_{f, g, Q_c + \phi_c?})$ is pairwise consistent. We know it contains exactly one complete answer s_n . In fact, any partial answer p in the set must be consistent with s_n . Otherwise, by the deontic constraint, there must be some complete answer distinct from s_n in $PBEST_n(w, f, g, S_{f, g, Q_c + \phi_c?})$ which entails p . But we just proved that cannot happen.

By definition, $S_{Q_c + \phi_c?, f, g}$ meets condition 2. To see that it meets condition 3, suppose that $Q' < S_{Q_c + \phi_c?, f, g}$ and Q' meets conditions 1 and 2. Then for some $s_n \in S_{Q_c + \phi_c?, f, g}$, Q' must contain complete answers q, q' such that $q \cup q' = s_n$. By a similar argument to that for condition 1, we can show that q and q' must be n th best relative to Q' . But since q and q' are complete answers, the n th best answers to Q' are not consistent after all.

References

- Beaver, David I. & Brady Z. Clark. 2008. *Sense and sensitivity: How focus determines meaning*. Oxford: Blackwell. <https://doi.org/10.1002/9781444304176>.
- Boylan, David & Ginger Schultheis. 2021. How strong is a counterfactual? *Journal of Philosophy* 118(7). 373-404. <https://doi.org/10.5840/jphil2021118728>.
- Brink, David O. 1994. Moral conflict and its structure. *Philosophical Review* 103(2). 215-247. <https://doi.org/10.2307/2185737>.
- Buchak, Lara. 2013. *Risk and rationality*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199672165.001.0001>.
- Cariani, Fabrizio. 2013a. 'Ought' and resolution semantics. *Noûs* 47(3). 534-558. <https://doi.org/10.1111/j.1468-0068.2011.00839.x>.
- Cariani, Fabrizio. 2013b. Epistemic and deontic *should*. *Thought: A Journal of Philosophy* 2(1). 73-84. <https://doi.org/10.1002/tht3.61>.
- Cariani, Fabrizio. 2016a. Consequence and contrast in deontic semantics. *Journal of Philosophy* 113(8). 396-416. <https://doi.org/10.5840/jphil2016113826>.

- Cariani, Fabrizio. 2016b. Deontic modals and probabilities: One theory to rule them all? In Nate Charlow & Matthew Chrisman (eds.), *Deontic modality*, 11–46. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198717928.003.0002>.
- Carr, Jennifer. 2015. Subjective *ought*. *Ergo* 2(27). 678–710. <https://doi.org/10.3998/ergo.12405314.0002.027>.
- Carter, Sam & John Hawthorne. 2021. Normality. Unpublished Manuscript.
- Castaneda, Hector-Neri. 1981. The paradoxes of deontic logic: The simplest solution to all of them in one fell swoop. In Risto Hilpinen (ed.), *New studies in deontic logic: Norms, actions, and the foundations of ethics*. Dordrecht: D. Reidel. https://doi.org/10.1007/978-94-009-8484-4_2.
- Charlow, Nate. 2016. Decision theory: Yes! Truth conditions: no! In Nate Charlow & Matthew Chrisman (eds.), *Deontic modality*, 47–81. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198717928.003.0003>.
- Charlow, Nate & Matthew Chrisman. 2016. *Deontic modality*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198717928.003.0003>.
- Copley, Bridget. 2006. What should “should” mean? Manuscript presented at the workshop *Language under uncertainty*, Kyoto, January 2006. <https://shs.hal.science/halshs-00093569/document>.
- Finlay, Stephen. 2009. Oughts and ends. *Philosophical Studies* 143(3). 315–340. <https://doi.org/10.1007/s11098-008-9202-8>.
- von Fintel, Kai. 2012. The best we can (expect) to get? Challenges to the classic semantics for deontic modals. Paper presented in a session on Deontic Modals at the Central APA, February 17, 2012. <http://mit.edu/fintel/fintel-2012-apa-ought.pdf>.
- von Fintel, Kai & Sabine Iatridou. 2008. How to say ‘ought’ in foreign: The composition of weak necessity modals. In Jacqueline Guéron & Jacqueline Lecarme (eds.), *Time and modality*, 115–141. Dordrecht: Springer. https://doi.org/10.1007/978-1-4020-8354-9_6.
- van Fraassen, Bas. 1973. Values and the heart’s command. *Journal of Philosophy* 70(1). 5–19. <https://doi.org/10.2307/2024762>.
- Gillies, Thony. 2012. Default logic and ‘hafta’. Note from Central APA Author Meets Critics Session for John Horty’s *Reasons as defaults*. <https://philpapers.org/rec/GILDLA-2>.
- Goble, Lou. 1996. Utilitarian deontic logic. *Philosophical Studies* 82(3). 317–357. <https://doi.org/10.1007/BF00355312>.

- Goble, Lou. 2005. A logic for deontic dilemmas. *Journal of Applied Logic* 3(3-4). 461-483. <https://doi.org/10.1016/j.jal.2005.04.004>.
- Goble, Lou. 2013. Prima facie norms, normative conflicts, and dilemmas. In Dov Gabbay, John Horty, Xavier Parent, Ron van der Meyden & Leendert van der Torre (eds.), *Handbook of deontic logic and normative systems, vol. 1*. Milton Keynes: College Publications. <https://www.collegepublications.co.uk/downloads/handbooksooooo1.pdf>.
- Hare, Caspar. 2010. Take the sugar. *Analysis* 70(2). 237-247. <https://doi.org/10.1093/analys/anp174>.
- Hawthorne, John. 2002. Lewis, the lottery and the preface. *Analysis* 62(3). 242-251. <https://doi.org/10.1093/analys/62.3.242>.
- Heim, Irene & Angelika Kratzer. 1998. *Semantics in generative grammar*. Oxford: Blackwell.
- Hoek, Daniel. 2018. Conversational exculpature. *Philosophical Review* 127(2). 151-196. <https://doi.org/10.1215/00318108-4326594>.
- Horty, John. 2012. *Reasons as defaults*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199744077.001.0001>.
- Horty, John. 2014. Deontic modals: Why abandon the classical semantics? *Pacific Philosophical Quarterly* 95(4). 424-460. <https://doi.org/10.1111/papq.12044>.
- Jackson, Frank. 1985. On the semantics and logic of obligation. *Mind* 94(374). 177-195. <https://doi.org/10.1093/mind/XCIV.374.177>.
- Jeffrey, Richard C. 1965. *The logic of decision*. Chicago: University of Chicago Press.
- Joyce, James M. 1999. *The foundations of causal decision theory*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511498497>.
- Karttunen, Lauri. 1977. Syntax and semantics of questions. *Linguistics and Philosophy* 1(1). 3-44. <https://doi.org/10.1007/BF00351935>.
- Khoo, Justin. 2021. Disjunctive antecedent conditionals. *Synthese* 198. 7401-7430. <https://doi.org/10.1007/s11229-018-1877-6>.
- Kratzer, Angelika. 1977. What 'must' and 'can' must and can mean. *Linguistics and Philosophy* 1. 337-355. <https://doi.org/10.1007/BF00353453>.
- Lassiter, Daniel. 2011. *Measurement and modality: The scalar basis of modal semantics*. New York, NY: New York University dissertation. <https://semanticsarchive.net/Archive/WMzOWU2O/Lassiter-diss-Measurement-Modality.pdf>.

- Lewis, David. 1996. Elusive knowledge. *Australasian Journal of Philosophy* 74(4). 549–567. <https://doi.org/10.1080/00048409612347521>.
- Makinson, David C. 1965. The paradox of the preface. *Analysis* 25(6). 205–207. <https://doi.org/10.1093/analys/25.6.205>.
- Mandelkern, Matthew. 2019. Bounded modality. *Philosophical Review* 128(1). 1–61. <https://doi.org/10.1215/00318108-7213001>.
- Mandelkern, Matthew, Ginger Schultheis & David Boylan. 2017. Agentive modals. *Philosophical Review* 126(3). 301–343. <https://doi.org/10.1215/00318108-3878483>.
- McNamara, Paul. 2004. Agential obligation as non-agential personal obligation plus agency. *Journal of Applied Logic* 2(1). 117–152. <https://doi.org/10.1016/j.jal.2004.01.006>.
- Moss, Sarah. 2015. On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics* 8(1). 1–81. <https://doi.org/10.3765/sp.8.5>.
- Ninan, Dilip. 2005. Two puzzles about deontic necessity. In Jon Gajewski, Valentine Hacquard, Bernhard Nickel & Seth Yalcin (eds.), *MIT working papers in linguistics*, vol. 51. Cambridge, MA: MITWPL. <https://philpapers.org/archive/NINTPA.pdf>.
- Portner, Paul. 2009. *Modality*. Oxford: Oxford University Press.
- Roberts, Craige. 2012. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics* 5. 1–69. <https://doi.org/10.3765/sp.5.6>.
- Rooth, Mats. 1992. A theory of focus interpretation. *Natural Language Semantics* 1(1). 75–116. <https://doi.org/10.1007/BF02342617>.
- Ross, W.D. 1930. *The right and the good*. Oxford: Oxford University Press.
- Rubinstein, Aynat. 2012. *Roots of modality*. Amherst, MA: University of Massachusetts, Amherst dissertation. <https://doi.org/10.7275/3527519>.
- Rubinstein, Aynat. 2014. On necessity and comparison. *Pacific Philosophical Quarterly* 95(4). 512–554. <https://doi.org/10.1111/papq.12047>.
- Savage, Leonard. 1972. *The foundations of statistics*. New York: Dover.
- Schroeder, Mark. 2010. *Ought*, agents, and actions. *Philosophical Review* 119(3). 1–41. <https://doi.org/10.1215/00318108-2010-017>.
- Schwarzschild, Roger. 1999. GIVENness, AVOIDF and other constraints on the placement of accent. *Natural Language Semantics* 7(2). 141–177. <https://doi.org/10.1023/A:1008370902407>.
- Sinnott-Armstrong, Walter. 1985. Moral dilemmas and incomparability. *American Philosophical Quarterly* 22(4). 321–329. <http://www.jstor.org/stable/20014112>.

Putting *oughts* together

- Snedegar, Justin. 2012. Contrastive semantics for deontic modals. In Martijn Blaauw (ed.), *Contrastivism in philosophy: New perspectives*. London: Routledge.
- Stalnaker, Robert. 1981. Indexical belief. *Synthese* 49(1). 129–151. <https://doi.org/10.1007/BF01063917>.
- Stalnaker, Robert. 1986. Possible worlds and situations. *Journal of Philosophical Logic* 15(1). 109–123. <https://doi.org/10.1007/BF00250551>.
- Stalnaker, Robert. 2014. *Context*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199645169.001.0001>.
- Swanson, Eric. 2016. The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic* 45(2). 121–146. <https://doi.org/10.1007/s10992-015-9367-5>.
- Yalcin, Seth. 2016. Modalities of normality. In Nate Charlow & Matthew Chrisman (eds.), *Deontic modality*, 230–255. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198717928.003.0009>.

David Boylan
Department of Philosophy, Texas Tech
University
Lubbock, TX 79409
dboylan@ttu.edu