

CONSEQUENCE AND CONTRAST IN DEONTIC SEMANTICS.

Fabrizio Cariani
Northwestern University.

Forthcoming in *The Journal of Philosophy*

Abstract: Contrastivists view *ought*-sentences as expressing comparisons among alternatives. Deontic actualists believe that the value of each alternative in such a comparison is determined by what would actually happen if that alternative were to be the case. One of the arguments that motivates actualism is a challenge to the principle of agglomeration over conjunction—the principle according to which if you ought to run and you ought to jump, then you ought to run and jump. I argue that there is no way of developing the actualist insight into a logic that invalidates the agglomeration principle without also invalidating other desirable patterns of inference. After doing this, I extend the analysis to other contrastive views that challenge agglomeration in the way that the actualist does. This motivates skepticism about the actualist’s way of challenging agglomeration.

Acknowledgments: Thanks to Steve Finlay and Melissa Fusco for their comments at the 2016 Eastern APA; to Paolo Santorio, Justin Snedegar and an anonymous referee for *The Journal of Philosophy* for feedback on previous versions. Also thanks to David Braun, Janice Dowell, Simon Goldstein, Rory Harder, Jeff Horta, Daniel Lassiter, Ben Lennertz, John Mackay, Shyam Nair, Jacob Ross, Mark Schroeder, and Stephen White for conversations and exchanges.

CONSEQUENCE AND CONTRAST IN DEONTIC SEMANTICS.

Ordinary speakers and theorists share the insight that sentences involving *ought*, like “Carla ought to run”, have an implicitly comparative nature. More precisely, the insight is that these sentences express comparisons between alternatives: one ought to run just in case running is better than the alternatives. Developing this insight into a theory requires an account of how the value of an alternative is determined. A prominent account has it that the value of an alternative is determined by what would happen if that alternative were the case. Translated into possible world talk, running is better than cycling just in case the closest world (to the actual world) in which one runs is better, in some salient way, than the closest world (to the actual world) in which one cycles.¹ Frank Jackson calls this view ‘actualism’.²

In “On the Semantics and Logic of Obligation”, Jackson (1985) motivates actualism on the basis of challenges to classical inference patterns for *ought*. Specifically, he targets the closure of *ought* under entailment and a principle of agglomeration over conjunction:

Conjunctive Inheritance: $ought(A \ \& \ B) \models ought(A)$

¹ The view assumes that a relation of betterness among individual worlds is given as primitive. It also presupposes that it makes sense to refer to *the* closest world to the actual world. There is a familiar debate about the uniqueness assumption built into talk of “the closest world to the actual world.” Famously, the semantics in Lewis’s *Counterfactuals* (London: Blackwell, 1973) denies uniqueness, while Stalnaker’s presupposes it—see his “A theory of conditionals,” in *Studies in Logical Theory, American Philosophical Quarterly Monograph Series 2* (1968) 98–112. Jackson acknowledges the point on p. 185 of “On the Semantics and Logic of Obligation,” *Mind* 94 (1985), while also noting that the view can be modified to satisfy Lewisians. Indeed, Goble developed a version of this logic that does not presuppose uniqueness. See, Lou Goble “A logic of ‘good, should’ and ‘would’: Part I,” *Journal of Philosophical Logic* 19 (1990) 169–199, Lou Goble “A logic of ‘good, should’ and ‘would’: Part II,” *Journal of Philosophical Logic* 19 (1990) 253–276, Lou Goble “The Logic of Obligation, ‘better’ and ‘worse’,” *Philosophical Studies* 70 (1993) 133–163. Like Jackson, however, I can convey all of my main points by focusing on the simpler version of the view.

² The conceptual roots of actualism are developed in Holly Smith Goldman “Dated rightness and moral imperfection,” in *The Philosophical Review* 85 (1974) 449–487 and Jordan Howard Sobel “Utilitarianism and Past and Future Mistakes,” in *Noûs* 10 (1976) 195–219. To my knowledge, Jackson, *op. cit.*, was the first to think of actualism as a semantic theory for the objective *ought* and to pose the question about the logic resulting from this theory. See also Jackson and Pargetter “Ought, Options and Actualism,” in *The Philosophical Review* 95 (1986) 233–255.

Agglomeration: $ought(A), ought(B) \models ought(A \& B)$

Since the examples that motivate these challenges are controversial, the debate about actualism has focused on its motivation. As a result, the view is rarely developed and important conceptual gaps remain concerning how it ought to be developed.

This paper explores one such gap. My central question is: how do standard accounts of logical consequence apply to the actualist semantic theory? ³ In particular, if logical consequence is understood as preservation of truth at a context, how should the actualist understand ‘truth at a context’ for sentences involving *ought*?

Filling out this gap is an essential part of the actualist agenda: after all, if actualism is motivated by claims about the unacceptability of certain inference patterns, it needs to make unequivocal predictions about the acceptability (or unacceptability) of inferences. One of the central results of this paper is that there are (at least) two very different ways of developing the main actualist principles into a notion of truth at a context. These yield radically different logics for *ought*.

Ultimately, this analysis results in a challenge for actualists. I argue that there is no notion of consequence that is both theoretically satisfactory and compatible with the cases that motivate the actualist. In other words, any attempt to satisfy the actualist’s motivation (which I describe in §2) induces a logic for *ought* that is far weaker than the literature has acknowledged (this is the main result of §3). This is because the only logics that are fully compatible with the actualist’s motivation additionally invalidate a principle that ought to be valid, namely:

Weakening: $ought(A), ought(B) \models ought(A \vee B)$

The news for actualism is not all bad: I will show (§4) that there is a viable version of actualism that validates Weakening. However, it is only available if the actualist is willing to give up the challenge to Agglomeration.

Having established these claims, I generalize the analysis beyond actualism. I consider two accounts of deontic *ought* that accept the

³ I am not the first to pose this question. It is the subject of the important series of papers by Lou Goble I referenced in footnote 1 above. My analysis will emphasize different aspects of the subject matter than those illuminated by Goble’s.

contrastivist idea (the idea that *ought*-sentences express comparisons among alternatives), but not the actualist's method for grounding such comparisons. The first account compares alternatives on the basis of their expected value. I show (§5) that this account shares the logical defects of the actualist semantics (it ends up invalidating Weakening). The second account, recently defended by Stephen Finlay, compares alternatives on the basis of their probability given that some salient end is realized. This account manages to invalidate Agglomeration while also validating Weakening. I will argue, however, that it lacks clear motivation for invalidating Agglomeration and cannot borrow the cases that motivate the actualist (§6).

I BACKGROUND: TYPES OF OUGHT-SENTENCES

This section and the next offer essential background. They may be skimmed by those who are familiar with the debate.

It is traditional to draw a distinction between two kinds of *ought*-sentences: those that specify what ought to be the case and those that specify what an agent (at a point in time, with a given body of information) ought to do.⁴ This distinction is not explicitly marked by the surface grammar of English. The sentence,

(1) Joan ought to ski at Mt. Bachelor.

is potentially ambiguous between a claim about what the world ought to be like (regardless of any agent's deeds) and one that states what Joan ought to do. Following Jackson, I will primarily, though not exclusively, focus on *ought-to-be* interpretations.

There is a second traditional distinction between subjective and objective *ought*. Roughly, one can distinguish between what Joan ought to do given her information (the 'subjective' reading) and what she ought to do given what in fact is the case (the 'objective' reading). These can come apart, for example, if Joan has misleading or insufficient information. Languages like English seem to allow both kinds of *ought-to-do* sentences. In principle, though this is less often discussed, there should also be a parallel subjective/objective distinction for *ought-to-be*. After all, there is a clear distinction between what ought to be the case

⁴ For one way of drawing this distinction, see Lloyd Humberstone "Two Sorts of Ought", in *Analysis* 32 (1971) 8-11.

in light of someone's (e.g. the speaker's) information and what ought to be the case in light of some body of facts. Actualism is meant to be a theory of the objective *ought*.

Finally, it is widely accepted that *ought*-sentences are context sensitive in a special way. Even if one focuses on a particular reading of (1), say the subjective *ought-to-do* reading, there is still no fixed, context-independent source of value to ground the *ought*-judgment. In some contexts, a speaker who uttered (1) may intend to convey that it is true because Joan's family will be at Mt. Bachelor and Joan wants to be with them. In other contexts, the speaker might mean that she needs to be in Eugene next week and wants to ski over the week-end. Kratzer famously interprets phenomena like this as instances of context dependence.⁵ The different interpretations arise from the value of a contextual parameter that records the relative priority of worlds (she calls this the 'ordering source'). Some of the views I consider here do not agree with the use of ordering sources, but they all accept the underlying indeterminacy in *ought*-sentences.

II BACKGROUND: DEONTIC ACTUALISM

Deontic actualism is a kind of contrastivism. Contrastivism is the thesis that *ought*-sentences have their truth-conditions relative to a class of alternatives. That is, the semantic value of $\ulcorner \textit{ought } A \urcorner$ is only fixed relative to a class of alternatives to the proposition that *A*. I use 'the proposition that *A*' as a term of art for 'the set of worlds in which *A* is true'. In the following, I use bold letters **A**, **B**, **C** to denote the set of worlds at which sentences *A*, *B*, *C* (respectively) are true. Though I use distinct styles of variables for these objects, I use sentential operators instead of set-theoretic notation (so for instance, I will write $\sim \mathbf{A}$ instead of $\overline{\mathbf{A}}$). With these conventions in place, the alternatives to **A** are non-empty mutually exclusive propositions (some authors require the alternatives to be exhaustive: this will not be necessary here).

The informal insight that motivates contrastivists is that competent speakers of English accept $\ulcorner \textit{It ought to be that } A \urcorner$ exactly when **A** is better than the alternatives.

Actualism is a thesis about what grounds such comparisons. I

⁵ See Angelika Kratzer "The Notional Category of Modality" in H. J. Eikmeyer & H. Rieser, eds., *Words, worlds, and contexts: New approaches in word semantics* (Berlin: de Gruyter, 1981) 38–74 and Angelika Kratzer, *Modal and Conditionals* (Oxford & New York: Oxford University Press, 2012).

sketched the view in the introduction, but it is important to be a bit more precise. Let us start by getting clear about the ingredients of the actualist semantics. First, for every **A**, there must be a unique world w such that w is the closest **A**-world to the actual world (w is an **A**-world iff $w \in \mathbf{A}$). When this relationship holds, I say that w is the *representative* of **A** (relative to the actual world). The second ingredient is a partial order $>$ defined on individual worlds.⁶ This partial order captures the relative goodness of worlds. Finally, say that $\ulcorner \textit{ought } A \urcorner$ is true relative to $\{\mathbf{A}, \mathbf{A}_1, \dots, \mathbf{A}_k\}$ just in case the representative of **A** is better (in the sense of $>$) than the representative of any other alternative.

As noted, Jackson motivates actualism on the basis of counterexamples to basic inference patterns involving *ought*, specifically:

Conjunctive Inheritance: $\textit{ought}(A \ \& \ B) \models \textit{ought}(A)$

Agglomeration: $\textit{ought}(A), \textit{ought}(B) \models \textit{ought}(A \ \& \ B)$

It is useful to separate two elements of Jackson’s challenge: the *empirical* claim is that these principles have intuitive counterexamples. The *theoretical* claim is that the actualist semantics predicts that they fail.

My focus is squarely on the theoretical claim—specifically, on what it takes to make it true. Nonetheless, some remarks on the empirical claim are in order. The attack on Inheritance is based on cases like Professor Procrastinate—who ought to accept a commitment to write a book review and write the review, but ought not to accept the commitment, for if he accepts, he will not write.⁷ Procrastinate cases are controversial and extensively discussed (again, I am setting aside whether they are convincing).⁸

⁶ A partial order is a transitive, reflexive and anti-symmetric relation. One could assume that $>$ has more structure than this, but that will not be necessary.

⁷ Jackson, “On the Semantics and Logic of Obligation”, p. 193

⁸ Even just focusing on very recent literature, variants of Prof. Procrastinate are discussed in Fabrizio Cariani “*Ought* and Resolutions Semantics,” *Noûs* 47 (2013) 534-558, Justin Snedegar “Contrastive Semantics for Deontic Modals,” in Martin Blauw, ed., *Contrastivism in Philosophy* (London: Routledge, 2012), Daniel Lassiter *Measurement and Modality: The Scalar Basis of Modal Semantics*, PhD Thesis, NYU (2012), Daniel Lassiter “Modality, Scale Structure, and Scalar Reasoning,” *Pacific Philosophical Quarterly*, 94 4 (2014) 461-490, Kai von Fintel “The best we can (expect to) get? Challenges to the classical semantics for deontic modals,” (2012) available at <http://mit.edu/fintel/fintel-2012-apa-ought.pdf>, Aaron Bronfman and Janice Dowell “Contextualism about Deontic Modals,” in Nate Charlow and Matthew Chrisman,

As for counterexamples to Agglomeration, Jackson adapts a variation on a case from Goldman (*op. cit.* §IX).⁹ Here is Jackson's description of the case:

Swerving 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. *It ought to be that Attila swerves, for then there would be no collision. [...] Equally it ought to be that Genghis swerves. But it ought not to be that both swerve, for then we get a worse collision.*¹⁰

I italicized Jackson's interpretation of the case. For a variety of reasons, I am skeptical that this is the correct diagnosis. There is room to diagnose the counterexample as turning on a slight equivocation: even if all three *ought* claims are interpreted as objective *ought-to-be* claims, they need not trade on the same contextual sources of value.¹¹ I will not pursue this criticism here, however. My response to Swerving Chariots

eds., *Deontic Modality* (Oxford: Oxford University Press, forthcoming), and Frank Jackson "Procrastinate Revisited," *Pacific Philosophical Quarterly* 95 (2014) 634-647.

⁹ If Jackson is right about Procrastinate cases, they would also be failures of Agglomeration.

(i) Procrastinate ought to accept and write.

(ii) Procrastinate ought not to accept.

(iii) Procrastinate ought to (accept and write and not to accept.)

(i) and (ii) sound jointly acceptable. By contrast, (iii) is contradictory. The counterexample I discuss in the main text purports to show something stronger: that agglomeration can fail even when the prejacent of the conclusion is consistent.

¹⁰ Jackson "On the Semantics and Logic of Obligation", p.189.

¹¹ The context sensitivity I noted in §1 is relevant to develop this alternative diagnosis. See also the discussion of similar cases in §5.4.3 of John F. Horty *Agency and Deontic Logic* (NY & Oxford: Oxford University Press, 2001), and §11.5.2 of his "Perspectival Act Utilitarianism," in Patrick Girard, Mathieu Marion and Olivier Roy, eds., *Dynamic Formal Epistemology* (Dordrecht: Springer, Synthese Library vol 351, 2011) 197-222.

will be indirect: even assuming that Jackson’s diagnosis is correct, it can be challenged on the basis of its consequences.

Let us move, then, to the theoretical claim. To vindicate this, an actualist must characterize a logical consequence relation that fits the semantics. The point should be obvious: to establish that actualism invalidates Agglomeration and Inheritance, one needs to know what ‘ \models ’ means. As I explain in the next two sections, what I have said so far about actualism is compatible with very different consequence relations.

Before presenting these consequence relations, let me comment on some additional logical desiderata that an actualist account should meet. Jackson notes that actualists might want to validate:

Negation Extraction: $ought(\sim A) \models \sim ought(A)$ ¹²

Incidentally, this emphasizes that the actualist’s rejection of Agglomeration is not based on the possibility of deontic conflicts. In addition, actualists endorse some multi-modal inferences—at least trivial sounding ones:

Intensional Invariance: $ought(A), necessarily(A \text{ iff } B) \models ought(B)$

Possibility: $ought(A) \models possibly(A)$

The last desirable inference pattern I highlight will be crucial to my argument:

Weakening: $ought(A), ought(B) \models ought(A \vee B)$

It is natural to view Weakening as a weakening of Agglomeration. In fact, if $\lceil ought(A \& B) \rceil$ entails $\lceil ought(A \vee B) \rceil$ and the consequence relation is transitive, Weakening follows from Agglomeration.¹³ Moreover, if the consequence relation is monotonic (that is, if adding premises preserves consequence), Weakening is entailed by (but does not entail) a general form of Inheritance—the claim that *ought* is closed under necessary entailment. Of course, actualists reject both Agglomeration and Inheritance, but that is compatible with accepting

¹² Jackson “On the Semantics and Logic of Obligation”, p.187.

¹³ Actualists will not be moved by this because they deny that $\lceil ought(A \& B) \rceil$ entails $\lceil ought(A \vee B) \rceil$.

Weakening. The next section establishes that, given the most natural characterization of the consequence relation, actualist semantics does not validate Weakening.

III LOGICAL CONSEQUENCE

How should an actualist define logical consequence? A very simple approach starts with the Kaplanian idea that logical consequence for context-dependent expressions is preservation of truth at a context.¹⁴ This account is appropriate in this theoretical setting because actualism treats *ought* as a context dependent expression: the semantic value of *ought* depends on several features of context: which world is actual, what is the closeness relation and which worlds are better than which others.

Definition 1 (Logical Consequence) $A_1, \dots, A_k \models B$ iff there is no context c such that all of the A_i are true at c but B is false at c .¹⁵

Given this definition, the actualist's task is to characterize truth at a context for *ought*-sentences. In turn, this requires her to say more about how the parameters she uses in her compositional semantics depend on context.

A simple model of context views each context as a quadruple of the form $\langle @, >, s, Alt \rangle$. In each quadruple, $@$ is the actual world, $>$ is the partial order that models relative goodness of worlds, s is a selection function that, given as inputs a proposition A and a world w , outputs the closest A -world to the actual world, and finally Alt is a set of alternatives (that is, a set of mutually exclusive non-empty propositions).

This model of context entails that, within a fixed context, no two sentences can be evaluated against different alternatives, and, for this

¹⁴ See pp. 522-523 of David Kaplan "On Demonstratives," in Joseph Almog, John Perry and Howard Wettstein, eds., *Themes from Kaplan* (New York & Oxford: Oxford University Press, 1989). One might, alternatively, formulate the actualist semantic theory as a Kripke-semantics for a modal logic with a deontic operator. Given such a formulation, one can leverage a standard model-theoretic analysis of consequence (e.g. preservation of validity in a frame). This approach has the advantage of allowing the use of standard techniques to study the metatheoretic properties of the semantics. The disadvantage is that it requires quite a bit of technical setup I will not need.

¹⁵ I have simplified Kaplan's definition by avoiding quantification over structures. None of my results are affected by this layer of complexity.

reason, it does not work well with actualism. Actualism requires:

Prejacent Representation: the alternative set against which $\ulcorner \textit{ought } A \urcorner$ is evaluated must include the set of worlds at which A is true.

Following established terminology, the *prejacent* of $\ulcorner \textit{ought } A \urcorner$ is the sentence A . It follows from Prejacent Representation that there is no set of alternatives (and hence, on the present view, no context) that allows simultaneous evaluation of:

- (2) It ought to be that Attila swerves
- (3) It ought to be that Genghis swerves
- (4) It ought to be that (Attila swerves and Genghis swerves).

There is no alternative set that includes the prejacent of (2), (3) and (4). Here is why: if the prejacent of (2) and (3) are not understood as incompatible with each other, then they cannot belong to the same alternative set. If they are, then their conjunction is a contradiction. If so, the prejacent of (4) does not belong to the alternative set (since every member of the alternative set is a non-empty proposition). This reasoning generalizes to every non-trivial instance of Agglomeration.¹⁶

To avoid this problem, actualists must abandon the claim that context alone determines a set of alternatives. Instead, they must hold that context *plus* the prejacent of a given *ought*-sentence determines the set of alternatives against which that particular sentence is evaluated.¹⁷

¹⁶ Justin Snedegar argues persuasively that any contrastivist must recognize the existence of cross-contextual reasoning, and with it some kind of cross-contextual account of acceptability of inferences. While I agree with Snedegar's claim, I reject the further claim (which to my knowledge Snedegar does not make) that this cross-contextual notion is the only notion involved in accounting for the validity and invalidity of inferences involving *ought*. See, Justin Snedegar "Deontic Reasoning Across Contexts," in Fabrizio Cariani, Davide Grossi, Joke Meheus and Xavier Parent, eds., *Deontic Logic and Normative Systems. Proceedings of DEON 2014* (Dordrecht: Springer, 2014).

¹⁷ See §2.2 of Wes Holliday's "Fallibilism and Multiple Paths To Knowledge", in Tamar Szabo Gendler and John Hawthorne, *Oxford Studies in Epistemology* vol. 5 (NY & Oxford: Oxford University Press, 2015) 97-144, for a similar distinction in the context of relevant alternatives accounts of knowledge. Holliday distinguishes theories that make the set of relevant alternatives dependent on the choice of proposition, as in Fred Dretske's "The Pragmatic Dimension of Knowledge," in *Philosophical Studies* 40, 3 (1981) 363-378, from theories that view the set of relevant alternatives as determined directly by context, as in Mark Heller's "Relevant Alternatives and

In other words, each context may ‘host’ different alternative sets (to be used depending on the prejacent of the given *ought*-sentence). Given this approach, it is possible to evaluate all of (2), (3) and (4) within a single context c . To formally implement this idea, model contexts as quadruples $\langle @, >, s, m \rangle$ with m a function from prejacent to alternatives. Call this function the *alternatives mapping*.

Not every function from prejacent to alternatives can count as an alternatives mapping, or else almost every argument will be invalid. Consider, for instance, the Negation Extraction pattern: $ought(\sim A) \models \sim ought(A)$ (recall that Jackson thinks actualism validates this). If the premise is evaluated relative to $\{\sim A, B\}$ and the *ought*-sentence in the conclusion is evaluated relative to $\{A, C\}$ for some unrelated C , the argument might fail. After all, the representative of A could be better than the representative of C even if the representative of $\sim A$ was better than the representative of B . The validity of Negation Extraction requires that the alternatives to A and the alternatives to $\sim A$ somehow be related.

Actualists need constraints on the alternatives mapping even for the inference patterns that are intended to fail (like Agglomeration). The failure of these patterns should not be grounded only in how alternatives are determined. The actualist semantics should play some role in explaining why they fail.¹⁸

Jackson is very clear about which alternatives mapping he favors: “I think that typically when we assert that it ought to be that A , we mean that it ought to be that A out of A and its negation”.¹⁹ Let’s give this a name:

Simple Mapping: map the prejacent A to the alternative set $\{A, \sim A\}$.

The simple mapping can be justified on grounds of simplicity (it is the simplest way of explaining how contrastive alternatives are generated) and generality (it applies to any prejacent without *ad hoc* stipulations).²⁰

Closure,” in *Australasian Journal of Philosophy* 77, 2 (1999) 197-208.

¹⁸ Without constraints on the alternatives mapping, Agglomeration would fail even for the actualist’s main opponents—the possibilists who maintain that the representative of A is the best possible A -world (regardless of closeness).

¹⁹ Jackson “On the Semantics and Logic of Obligation”, p.187.

²⁰ As defined in the main text, the simple mapping is context invariant. It can be made context dependent by adding another parameter to our representation of context. Let

Putting these ideas together yields a theoretical package consisting of four items: the actualist semantics, the thesis that logical consequence is preservation of truth at a context, the thesis that (fixing a context) we can associate different alternative sets to different prejacent, and finally the simple mapping.

It is easy to check that this package yields the verdicts that motivate the actualist. Inheritance and Agglomeration are invalid. Negation extraction is valid, as are the basic multi-modal patterns (Intensional Invariance and Possibility).

However, and crucially, Weakening fails. That is, there are contexts in which $\ulcorner \text{ought}(A) \urcorner$ and $\ulcorner \text{ought}(B) \urcorner$ are both true but $\ulcorner \text{ought}(A \vee B) \urcorner$ is false.

Counterexample to Weakening: The following assignments are sufficient to fix a context in which the premises are true and the conclusion is false (in the Closeness column "0" labels the actual world).

Possibility	Representative	Closeness	Goodness
A&B	$w_{A,B}$	1	best
A&~B	$w_{A,\sim B}$	2	worst
~A&B	$w_{\sim A,B}$	0	tied for second
~A&~B	$w_{\sim A,\sim B}$	3	tied for second

In this context, $\ulcorner \text{ought}(A) \urcorner$ is true because $w_{A,B}$, which represents **A**, is the best world; $\ulcorner \text{ought}(B) \urcorner$ is true because the actual world, which represents **B**, is better than the representative of $\sim \mathbf{B}$ (which is $w_{A,\sim B}$). However, $\ulcorner \text{ought}(A \vee B) \urcorner$ is false, because the representative of the disjunction is the actual world and the representative of its negation is $w_{\sim A,\sim B}$ and these worlds are equally good.²¹

f_c be a set of worlds consisting of all the worlds that are relevant in c .

Contextualist Simple Mapping: For every context c , for each prejacent A in context c , let the alternative set be $\{f_c \cap \mathbf{A}, f_c \cap \sim \mathbf{A}\}$

None of my findings depend on the choice between these two implementations of the simple mapping. For simplicity, I always discuss the context invariant one.

²¹ An anonymous referee and Ben Lennertz independently asked an interesting question about this counterexample: given that the actual world verifies $\sim A \& B$, shouldn't the closest $\sim \mathbf{B}$ -world be $w_{\sim A,\sim B}$ rather than $w_{A,\sim B}$. After all, both $w_{\sim A,\sim B}$ and the actual world verify $\sim A$. In response, I highlight two points: first, it is plausible that similarity between worlds depends on global features of the given worlds more so than on local features. That is to say: how matters are beyond **A** and **B**, might be more


The invalidity of Weakening is a serious problem for the simple actualist package I put together. Though the actualist’s motivation might warrant a non-normal modal logic, it does not warrant a logic this weak. Weakening is an eminently plausible principle and no example has been offered in the literature to justify its failure. Note, incidentally, that the intuitions in *Swerving Chariots* do not support the failure of Weakening. The sentence “It ought to be that (either Genghis or Attila swerves)” is generally judged true in the context. Moreover, the actualist semantics predicts this sentence to be true, so the intuitive pull against Weakening cannot possibly come from *Swerving Chariots*.

In response, the actualist must either bite the bullet and reject Weakening (in addition to Agglomeration) or avoid the package I set up. In the next section, I consider a way of modifying the theoretical package that can restore the validity of Weakening (but at a cost).

To further motivate this alternative package, it is useful to consider a second problem for the current version of actualism. The problem is that, by lumping together the alternatives to a prejacent into a single alternative, the simple mapping forces some mistaken predictions.²²

Harry’s Mischief. Suppose that Harry has three options. In order of goodness of their representative world, he can (S) go to school or (V) play videogames at home or (L) shoplift at the nearby mall. Suppose also that in the actual world he shoplifts at the mall. Suppose also that the representative of V is closer to the actual world than the representative of S.

A diagrammatic representation of the example helps illustrate some of the key points. The three depicted worlds are the representatives of




important for establishing similarity relations. Moreover, even at the local level, $w_{A, \sim B}$ and $w_{\sim A, B}$ are alike in that they both verify $A \equiv \sim B$, which is false in $w_{\sim A, \sim B}$. The closeness ordering in the counterexample might arise if $A \equiv \sim B$ were to be an important dimension of closeness (more important than $\sim A$).

²² This style of objection is independently due to Jacob Ross “Actualism, Possibilism and Beyond”, *Oxford Studies in Normative Ethics* 2 (2013) 74-96 and to a 2009 blog post by Ralph Wedgwood titled “Against Actualism” (available at <http://peasoup.typepad.com/peasoup/2009/09/against-actualism.html>).

each alternative. The actual world is depicted as a circle. Spatial proximity represents closeness to the actual world. Relative goodness is represented by $<$ (higher is better).

The current implementation of actualism implies counter-intuitively that Harry ought to play videogames at home, because the representative of V is better than the representative of $\sim V$.



Represents $\sim V$

The problem is that the good S -world is, so to speak, hidden behind the non-ideal representative of $\sim V$ which happens to be actual. Incidentally, the actualist package also implies that Harry ought to go to school, the representative of this alternative is better than the representative of its negation. The upshot is that actualism views Harry's Mischief as a situation in which two incompatible things ought to be. Maybe some deontic conflicts are like this, but it is implausible that Harry's Mischief is a deontic conflict.

In fact, the above reasoning shows that things are a bit worse than I've suggested: according to the current actualist package, one ought to do each of the things that are better than what one in fact does. If one will, in fact, lethally poison 100 cats, then for each $n < 100$, one ought to lethally poison exactly n cats.

The problem has a clear source: what matters to the truth of $\ulcorner \text{ought}(A) \urcorner$ is how A compares with each individual alternative. By lumping together the alternatives that are incompatible with A , one might lose sight of an alternative that undermines the claim that A ought to be the case (in Harry's case: going to school). I am going to call this *the lumping problem*.

In the next section, I discuss whether the lumping problem and the invalidity of Weakening can be avoided by moving away from the simple mapping.

IV ACTUALISM AND THE FINE MAPPING

Actualists can solve the lumping problem by contrasting the representative of A with the representative of each individual alternative to A .²³ According to the resulting view, $\ulcorner \text{ought}(A) \urcorner$ is true

²³ Justin Snedegar independently identified a version of this response to the lumping

in context c iff the representative of A is better than the representative of each alternative to A in c .

This idea allows an actualist-friendly account of Harry's Mischief: it is not the case that Harry ought to stay at home and play videogames. After all, the representative of V is not better than the representative of S (which is an alternative in its own right). Moving away from the simple mapping requires us to say something about how alternatives are generated and distinguished in context (this will result in a more complex package).

Start with the idea that each context determines a background set of alternatives. Very roughly, the background alternatives in c are all the possibilities that are worth distinguishing in c . Depending on which reading of *ought* is salient, 'worth distinguishing' may mean something slightly different. Since it is up to the actualist to tell a full story here, I will not try to say more. I will, however, offer some examples for determinacy's sake. In *Swerving Chariots*, the background alternatives must distinguish between the worlds in which Attila swerves and the worlds in which he does not (and similarly for Genghis). By contrast, in that context, the background alternatives need not distinguish the worlds in which Attila swerves while yelling "yee-hah!" from worlds in which he swerves without yelling at all. Similarly, in *Harry's Mischief*, the background alternatives should distinguish the three possible courses of actions available to Harry, but, for the purposes of the conversation, they need not distinguish finer possibilities.

There is no guarantee that the background alternatives can be deployed in evaluating a given *ought*-sentence. This is because the background alternatives need not satisfy Prejacent Representation. However, given an *ought*-sentence with prejacent A , and an appropriate set \mathcal{B} of background alternatives, it is possible to transform \mathcal{B} into a new set \mathcal{B}^A that does satisfy Prejacent Representation. Model each context as a quadruple: $\langle @, >, s, \mathcal{B} \rangle$ where \mathcal{B} is a background set of alternatives. Then claim that $\ulcorner \text{ought}(A) \urcorner$ in context c with background alternative set \mathcal{B}

- presupposes that A is a union of members of \mathcal{B} .
- Is evaluated against:

$$\mathcal{B}^A = \{A\} \cup \{B \in \mathcal{B} \mid B \text{ is incompatible with } A\}$$

problem and spelled it out in his unpublished manuscript "Ought and Options".

In algorithm form: check that \mathbf{A} is a union of alternatives in \mathcal{B} ; if it is, collect the alternatives compatible with \mathbf{A} into a single alternative, and then compare that single alternative against all the remaining alternatives. Call this the *fine mapping*.

Because the present account introduces presuppositions as part of the meaning of *ought*, the notion of logical consequence requires a tweak. We need something like what von Fintel (1999) calls ‘Strawson Entailment.’²⁴

Definition 2 $A_1, \dots, A_k \models_s B$ iff there is no context c such that (i) all of the presuppositions of each of the A_i ’s and of the conclusion B are satisfied in c (ii) all of the A_i are true at c but (iii) B is false at c .

Our new theoretical package then is this: actualism, the modified notion of logical consequence, the modeling of contexts as quadruples of the form $\langle @, >, s, \mathcal{B} \rangle$, where \mathcal{B} is a background set of alternatives, and finally the recipe for turning \mathcal{B} into \mathcal{B}^A .

There is much that is of interest in this package. It invalidates Inheritance for the reason Jackson wants. It validates Negation Extraction, and most surprisingly it validates Weakening. The lumping problem is also avoided (by design): it is false that Harry ought to play videogames at home, since going to school is better.

The problem with this package is that it is not faithful to the actualist’s motivation. Somewhat surprisingly, it validates Agglomeration. Proving this requires working through details that are best relegated to a note.²⁵ However, we can observe at a reasonably non-

²⁴ Kai von Fintel “NPI Licensing, Strawson Entailment and Context Dependency,” *Journal of Semantics*, 16, 2(1999): 97-148. The name is due to remarks by Strawson in “On Referring,” in *Mind* 59 235 (1950) 320-344. Strawson discusses notion of logical consequence that is appropriate to a setting in which proper names carry existence presuppositions. I have adapted Strawson entailment to the idea of logical consequence as preservation of truth at a context.

²⁵ Suppose for a contradiction that there is a context c such that (i) $\ulcorner \text{ought}(A) \urcorner$ is true in c , (ii) $\ulcorner \text{ought}(B) \urcorner$ is true in c and (iii) $\ulcorner \text{ought}(A\&B) \urcorner$ is false in c . Suppose also that \mathbf{A} , \mathbf{B} and $\mathbf{A\&B}$ are all unions of alternatives in the background alternative set of c , so that all the relevant presuppositions are satisfied. Because of (i), the representative of \mathbf{A} , w_A , is better than the representatives of each alternative that is incompatible with \mathbf{A} . Because of (ii), the representative of \mathbf{B} , w_B , is better than the representatives of each alternative that is incompatible with \mathbf{B} . We consider two cases: in case 1, w_A is the closest $\mathbf{A\&B}$ -world. In case 2, w_A is the closest $\mathbf{A\&\sim B}$ -world. These cases are exhaustive because we assume that the closest \mathbf{A} -world is either the closest $\mathbf{A\&B}$ -

technical level that the current package does not predict that Agglomeration should fail in Swerving Chariots. The premises of Swerving Chariots are now both false in the given context: the representative of Attila's swerving is not better than the representatives of each of the other alternatives, because it is not better than the representative of the worlds in which Attila does not swerve but Genghis does. This falsifies (2). Parallel reasoning establishes the falsity of (3).

This package forces significant revisions of the actualist's claims about logic. Actualists who are willing to give up the attack on Agglomeration can adopt the fine mapping and avoid both of the problems discussed in §3. However, my criticism of the actualist's challenge to Agglomeration stands: none of the packages I considered manages to invalidate Agglomeration while validating Weakening.

V EXPECTED VALUE CONTRASTIVISM INVALIDATES WEAKENING

The investigation can be extended to other contrastive theories. I consider two such theories. In this section, I consider views that combine contrastivism with the idea that alternatives ought to be

world or the closest **A&~B**-world.

Case 1: w_A is the closest **A&B**-world. Then w_A would also be the representative of **A&B**, which by (iii) is not better than all the incompatible alternatives. We know from (i) that w_A is better than the representatives of alternatives that entail $\sim \mathbf{A}$. It follows that the representative of $\sim (\mathbf{A}\&\mathbf{B})$, which is no worse than w_A , is the representative of some alternative that entails **A&~B**. Call this w' . We must have that it is not the case that $w_A > w'$. However we know that $w_B > w'$, because w_B is better than the representatives of each alternative to **B** (any alternative that entails **A&~B** is an alternative to **B**). Now w_B is either the closest **A&B**-world or the closest $\sim \mathbf{A} \& \mathbf{B}$ -world. In the first case, we have $w_A = w_B$, so that is a contradiction (since we know that $w_A > w'$ is false). In the second case, $w_B < w_A$ (by (i)) and that is also incompatible with the falsity of $w_A > w'$.

Case 2: w_A is the closest **A&~B**-world. Because of (ii), the representative of **B** (call it w_B) must be better than w_A (i.e. $w_B > w_A$). If w_B were the representative of $\sim \mathbf{A} \& \mathbf{B}$, then $\ulcorner \text{ought } A \urcorner$ would be false (because $w_B > w_A$). So w_B must be the representative of **A&B**. But then $\ulcorner \text{ought}(A\&B) \urcorner$ must be true, because w_B is better than w_A , and also better than the representatives of any alternative incompatible with **A** (because w_A must be better than any such alternative or else $\ulcorner \text{ought}(A) \urcorner$ would not be true). Each case is contradictory, so the assignment must be impossible.

compared on the basis of their expected values.²⁶ Such views are, in a sense, natural complements to actualism. One might view actualism as a contrastive account of the objective *ought*, and expectational theories as contrastive accounts of the subjective *ought*.²⁷

Given a numerical distribution of value over individual worlds, one can calculate the expected value of a set of worlds. The expected value of **A**, written $EV(\mathbf{A})$, is the average of the value of the worlds w that belong to **A** weighed by $Pr(w|\mathbf{A})$.²⁸ It is an elementary fact that expected values are not monotonic, in the sense that there are **A**, **B** such that $\mathbf{A} \subset \mathbf{B}$, but $EV(\mathbf{B}) \leq EV(\mathbf{A})$. Additionally, the expectation of a conjunction, $EV(\mathbf{A}\&\mathbf{B})$, can be significantly below both of $EV(\mathbf{A})$ and $EV(\mathbf{B})$.

A contrastive account of *ought* can be built on the notion of expected value: $\ulcorner \text{ought}(A) \urcorner$ is true in context c just in case the expected value of **A** is greater than the expected value of each of the other alternatives in c .²⁹ Once again, we can complete the theoretical package with the simple mapping or with the fine mapping. Let us start with

²⁶ See Lou Goble “Utilitarian Deontic Logic,” in *Philosophical Studies* 82 (1996) 317-357; Daniel Lassiter *Measurement and Modality: the Scalar Basis of Modal Semantics* PhD thesis, NYU (2011). I emphasize that the analysis in this section only applies to accounts that are contrastive in the sense that $\ulcorner \text{ought}(A) \urcorner$ is true iff **A** is better than the alternatives. There are expected value accounts that reject this form of contrastivism and do not fall prey to the objections I level in this section: see Ralph Wedgwood “Subjective and Objective Ought,” in N. Charlow & M. Chrisman, eds., *Deontic Modals* (New York & Oxford: Oxford University Press, forthcoming). I have formulated a broader critique of these views in Fabrizio Cariani “Deontic Modals and Probabilities: One Theory to Rule Them All?,” in Charlow and Chrisman, eds., *Deontic Modality*, *op. cit.*, pp. 11-46.

²⁷ Frank Jackson “Procrastinate Revisited,” in *Pacific Philosophical Quarterly* 95 4 (2014) 634-647.

²⁸ There are other ways of calculating expected values: this is how Goble, *op. cit.* does it. It suffices to make all the necessary points.

²⁹ Expectational contrastivists presuppose something that some decision theorists find philosophically dubious (even if mathematically trivial): namely that one can compute expected values of arbitrary conjunctions and disjunctions. For example, it presupposes that, in addition to the expected value of “Joe runs” and the expected value of “Joe jumps”, we can also calculate the expected value of “Joe runs or jumps”. This seems to presuppose that it is intelligible to talk about the prior probability of the free action of an agent even in the context of deliberation. For challenges to this idea, see Isaac Levi “Rationality, Prediction, and Autonomous Choice,” *Canadian Journal of Philosophy* 19 suppl. (1989) 339-363, Isaac Levi “Review Essay *The Foundations of Causal Decision Theory*,” in this JOURNAL, 97 374 (2000). For replies, see Wlodek Rabinowicz “Does Practical Deliberation Crowd Out Self-Prediction?,” *Erkenntnis* 57 (2002) 91-122 and James M. Joyce “Levi on Causal Decision Theory and the Possibility of Predicting One’s Own Actions,” in *Philosophical Studies*, 110 (2002) 69-102.

‘Simple Expectational Contrastivism’ (SEC for short). This is the view that $\ulcorner \text{ought}(A) \urcorner$ is true in context c iff $EV(\mathbf{A}) > EV(\sim \mathbf{A})$.³⁰

A simple model of context that fits this view treats each context as a quadruple of the form $\langle @, Pr, v, m \rangle$. In each quadruple, $@$ is the actual world, Pr a probability function, v a value function and m the simple mapping. (Further bells and whistles are usually added, but they are not necessary here.)

As far as our desiderata go, SEC behaves similarly to the simple version of actualism. It invalidates Agglomeration and Inheritance, while validating Negation Extraction and the multi-modal patterns.

It also runs into two familiar problems: first, it invalidates Weakening. There are contexts in which $EV(\mathbf{A}) > EV(\sim \mathbf{A})$, $EV(\mathbf{B}) > EV(\sim \mathbf{B})$, but $EV(\mathbf{A} \vee \mathbf{B}) \leq EV(\sim (\mathbf{A} \vee \mathbf{B}))$.

Counterexample to Weakening for SEC: There is a context $\langle @, Pr, v, m \rangle$ that makes the premises of Weakening true and its conclusion false. Consider Pr and v such that:

A	B	Pr	EV
T	T	.25	100
T	⊥	.25	-50
⊥	T	.25	-50
⊥	⊥	.25	0

In this context, $EV(\mathbf{A}) = 25$, while $EV(\sim \mathbf{A}) = -25$, making $\ulcorner \text{ought}(A) \urcorner$ true. Similarly, $EV(\mathbf{B}) = 25$, while $EV(\sim \mathbf{B}) = -25$, making $\ulcorner \text{ought}(B) \urcorner$ true. Finally, $EV(\mathbf{A} \vee \mathbf{B}) = 0 = EV(\sim (\mathbf{A} \vee \mathbf{B}))$, making $\ulcorner \text{ought}(A \vee B) \urcorner$ false.

SEC also shares the lumping problem with the simple version of actualism. By lumping together the alternatives to a prejacent into a single alternative, SEC predicts the truth of too many *ought*-sentences.

Harry’s Expectational Mischief. There are three possible things Harry might do: he can (S) go to school or (V) play videogames at home or (L) shoplift at the nearby mall.

³⁰ To an approximation, this is the account of Goble *op. cit.* and Lassiter *op. cit.*. Lassiter’s variant requires the expected value of \mathbf{A} to significantly exceed the expected value of $\sim \mathbf{A}$. This does not matter to any of the results to follow.

<i>alternative</i>	<i>Pr</i>	<i>EV</i>
S	1/3	10
V	1/3	-5
L	1/3	-50

The particular numbers do not matter very much. What matters is that *SEC* implies counterintuitively that Harry ought to play videogames at home, since **V** is better than $\sim \mathbf{V}$ (note: $EV(\sim \mathbf{V}) = .5 \cdot 10 + .5 \cdot -50 = -20$). In addition, like actualism, *SEC* also predicts that Harry ought to go to school—effectively, treating Harry’s case as a moral dilemma of sorts (which it clearly is not).

Harry’s Expectational Mischief shows that the expectational contrastivist is best served by adopting something like the fine mapping instead. According to the resulting view, $\ulcorner \text{ought}(A) \urcorner$ is true in c just in case the expected value of A is greater than the expected value of each of the alternatives. In §4, I noted that actualism plus the fine mapping validates Agglomeration and Weakening. This does not happen in the context of expected value contrastivism. Both patterns remain invalid even if we move to the fine mapping. The counterexample to Weakening is, in fact, a simple extension of the one that worked for *SEC*.³¹

The upshot is this: like actualists, expectational contrastivists cannot invalidate Agglomeration without also invalidating Weakening. But invalidating Weakening is implausible (at any rate, expectational contrastivists have never tried justifying it). This undermines their case against Agglomeration and raises the question: is it even possible to invalidate Agglomeration without invalidating Weakening? To answer this question, I consider one last kind of contrastive theory.

³¹ *Counterexample to Weakening*: Consider the context set out in counterexample against *SEC* and add the background alternative set,

$$\{\mathbf{A\&B}, \mathbf{A\&\sim B}, \sim \mathbf{A\&B}, \sim \mathbf{A\&\sim B}\}$$

Establishing $\ulcorner \text{ought}(A) \urcorner$, $\ulcorner \text{ought}(B) \urcorner$, $\ulcorner \sim(\text{ought}(A \vee B)) \urcorner$ requires a few new checks. Note that the presuppositions of all the *ought*-sentences are satisfied by the context. In the original example, $EV(\mathbf{A}) = 25$, which exceeds both $EV(\sim \mathbf{A} \& \mathbf{B})$ and $EV(\sim \mathbf{A} \& \sim \mathbf{B})$. Similarly, $EV(\mathbf{B}) = 25$ which exceeds both $EV(\mathbf{A} \& \sim \mathbf{B})$ and $EV(\sim \mathbf{A} \& \sim \mathbf{B})$. Finally, $EV(\mathbf{A} \vee \mathbf{B}) = 0$, which falsifies the conclusion, because $EV(\sim (\mathbf{A} \vee \mathbf{B})) = 0$ as well.

VI END-RELATIONAL CONTRASTIVISM

The last contrastive theory I consider is what Stephen Finlay calls the *end-relational* analysis.³² End-relationalism relativizes the truth-conditions of *ought*-sentences to a probability function Pr and an end E . There are a couple ways of spelling out this insight:

Inverse end-relationalism. $\lceil \textit{ought}(A) \rceil$ is true relative to $\{A, A_1, \dots, A_k\}$ iff $Pr(A|E) > Pr(A_i|E)$ for each $i \in \{1, \dots, k\}$

Direct end-relationalism. $\lceil \textit{ought}(A) \rceil$ is true relative to $\{A, A_1, \dots, A_k\}$ iff $Pr(E|A) > Pr(E|A_i)$ for each $i \in \{1, \dots, k\}$

Simplifying a bit, on the first version, what matters is the probability of each option given the relevant ends. On the second, what matters is the probability of the end given the relevant options. Finlay's official account of *ought* in a sense involves both versions. According to him, the truth-conditions of deontic *ought* are effectively approximated by Direct end-relationalism. However, the most general lexical entry for *ought*, regardless of interpretation, is the Inverse account (which Finlay deems more adequate for sentences like "It ought to rain"). To swap the direction of the conditional probabilities, he adds the assumption that in deontic interpretations every option must have the same prior probability (he calls this Symmetry of Choice).³³

Like the previous theories, end-relationalism is meant to make predictions about inferences, but rarely, if ever, accompanied by a definition of logical consequence. As in those cases, we hypothesize that logical consequence is preservation of truth at a context. Once again, then, we must get clear about what contexts look like for end-relationalists. At a minimum, contexts must contain, in addition to the

³² See Stephen Finlay, "Ought and Ends," in *Philosophical Studies* 143 3 (2009) 315-340; Stephen Finlay, *Confusion of Tongues* (New York & Oxford: Oxford University Press, 2014); Stephen Finlay "Ought: Out of Order" in N. Charlow & M. Chrisman, eds., *Deontic Modals* (New York & Oxford: Oxford University Press, forthcoming).

³³ For a much more thorough explanation, see Finlay *Confusion of Tongues*, *op cit.*, §3.6. I have criticized Symmetry of Choice in Fabrizio Cariani "Epistemic and Deontic *Should*," in *Thought* 2 1 (2013) 73-84. For Finlay's replies to some of my objections, see his "Ought: Out of Order", *op. cit.*.

actual world, a probability function Pr and an end E .³⁴ As above, we can construct two implementations of end-relationalism depending on whether we use the simple or the fine mapping. To keep the variety of theories to a minimum I will only talk about versions of the view that use the simple mapping.

At first sight, it might look that Inverse End-Relationalism might do better. First, it validates a general form of Inheritance.³⁵

Inheritance: if $A \subseteq B$, $ought(A) \models ought(B)$.

Given Inheritance, Weakening is a small step away. If the consequence relation is monotonic, one can trivially reason from a less general, disjunctive version of Inheritance to Weakening.

Disjunctive Inheritance: $ought(B) \models ought(A \vee B)$

Weakening: $ought(A), ought(B) \models ought(A \vee B)$

Since both our concepts of logical consequence (Definitions 1 and 2) are monotonic, Weakening follows.³⁶ Despite validating Weakening, Inverse End-Relationalism invalidates Agglomeration.³⁷

One might think, then, that Inverse End-Relationalism is a

³⁴ Strictly speaking, Finlay denies this. He prefers to say that contexts contain a qualitative information state I (e.g. a set of worlds) and that I determines a probability function Pr . This difference is irrelevant to the present argument.

³⁵ In the case of the simple mapping variant, this follows because if $A \subseteq B$ and $Pr(A | E) > Pr(\sim A | E)$, we must also have that $Pr(B | E) > Pr(\sim B | E)$. In the case of the fine mapping variant, if $Pr(A | E) > Pr(C | E)$ for each alternative C to A , we must also have that $Pr(B | E) > Pr(D | E)$ for each alternative D to B . The reason is that the alternatives to B (according to the fine mapping) form a subset of the alternatives to A and, by the probability calculus $Pr(B | E) \geq Pr(A | E)$.

³⁶ The assumption of monotonicity of the consequence relation has been challenged by researchers in the dynamic semantics tradition. In the context of deontic modality, see for instance Malte Willer's "A Note on *Iffy* Oughts," this JOURNAL 109 449-461 and Malte Willer "Dynamic Foundations for Deontic Logic," in Nate Charlow and Matthew Chrisman, eds., *Deontic Modality* (Oxford: Oxford University Press, forthcoming). However, the particular instance of monotonicity involved here is very unlike the sort of pattern that is challenged by dynamic semanticists. Moreover, end-relationalism is typically presented in a static framework.

³⁷ To see how it can fail, consider the following model: $Pr(A \& B) = .1$; $Pr(A \& \sim B) = .45$; $Pr(\sim A \& B) = .45$; $Pr(\sim A \& \sim B) = 0$. Then $Pr(A | E) > Pr(\sim A | E)$; $Pr(B | E) > Pr(\sim B | E)$, but $Pr(A \& B | E) > Pr(\sim(A \& B) | E)$. This establishes that Agglomeration is invalid for the simple mapping version of Inverse End-Relationalism. The same model, but with different inequalities, shows that it fails for the fine mapping version.

promising option for those who want to rejection Agglomeration because of cases like Swerving Chariots. This is not so: Inverse End-Relationalism does not predict that Swerving Chariots is a counterexample to Agglomeration. Suppose, for instance, that the end E is to avoid a collision. Now, if $Pr(\text{Attila Swerves}|E) = Pr(\text{Genghis Swerves}|E)$, then the premises of Swerving Chariots are false. The probability that a collision is avoided because Attila but no Genghis swerves equals the probability that a collision is avoided because Genghis swerves but Attila does not. On the other hand, if $Pr(\text{Attila Swerves}|E) > Pr(\text{Genghis Swerves}|E)$, it is false that Genghis ought to swerve, for after all given the end it is more likely that Attila swerves.

Let me reiterate that I find these to be reasonable predictions about Swerving Chariots. The objection against Inverse end-relationalism is not that it makes the wrong predictions about Swerving Chariots. It is, instead, that it violates Agglomeration but it cannot appeal to the standard counterexamples against it. In addition to this, the particular way in which Agglomeration fails for Inverse end-relationalists is not supported by intuitive cases.

As I mentioned, however, Finlay's official account of deontic *ought* amounts to Direct end-relationalism. As it turns out, this view is a special case of expectational contrastivism (so we can be relatively brief in bringing out its logical properties). To see this, suppose that the value of a world w is 1 if E is true at w and 0 otherwise. Then, for the given value function v and probability Pr : $EV(A) = Pr(A|E)$, where $EV(\bullet)$ is the expected value function determined by v and Pr . One might then expect that Direct end-relationalism and expectational contrastivism determine similar logics. This is exactly what turns out to be the case: specifically, both invalidate Agglomeration and Weakening (as well as all the versions of Inheritance I discussed in this paper). If, as I maintain, it is imperative to hold on to Weakening, then Direct end-relationalism is in as much trouble as the views discussed in the previous sections.

VII CONCLUSIONS

The central finding of this paper is that there are systematic problems in setting up a contrastive theory that (i) invalidates Agglomeration, (ii) validates Weakening and (iii) is faithful to the actualist's motivation for invalidating Agglomeration. Actualism, expectational contrastivism and Direct end-relationalism generally violate constraint (ii). I built a

version of Actualism that satisfies constraint (ii), but the cost is the failure of (i) and (iii). Inverse end-relationalism has the distinctive advantage of satisfying (i) and (ii), but it violates constraint (iii).

Interpreting this conclusion requires some care. I do not suggest that Agglomeration ought to be generally valid. Many theorists believe it should fail in the case of true deontic conflicts. It is remarkable here that the logics that are designed to model deontic conflicts do validate Weakening even as they invalidate Agglomeration.³⁸

Neither do I intend these conclusions to be an indictment of the contrastive insight that motivated the views I considered here. There are contrastive theories that validate Agglomeration and simply deny that Swerving Chariots is evidence against Agglomeration.³⁹ Furthermore, there may be other ways of characterizing the notion of logical consequence I have not considered that satisfy all three desiderata (though I think there is reason to doubt this).

The present discussion also holds a positive moral for contrastivist views more generally. Developing a contrastive view requires explicit choices about how the alternatives interact with context. In particular, contrastivists must be extremely careful in specifying what kind of alternatives mapping they choose. Even assuming that logical consequence is preservation of truth at a context, the logical properties of a contrastive account may vary significantly according to what role is assigned to context in determining the set of alternatives that is appropriate for a given *ought*-sentence.

³⁸ The suggestion that Agglomeration might be involved in moral dilemmas comes from the influential essay by Bernard Williams “Ethical Consistency,” *Proceedings of the Aristotelian Society* (Supplement), 39 (1965) 103-124. For an example of a logic that invalidates Agglomeration, but not Weakening, in the face of moral dilemmas, see the proposal of Bas van Fraassen “Values and the Heart’s Command,” this JOURNAL 70 (1973) 5-19.

³⁹ I have presented one such theory in Fabrizio Cariani “ ‘Ought’ and Resolution Semantics”, *op. cit.*