

A problem about preference*

Anthony S. Gillies
Rutgers University

1 Nutshell (introductory)

Obligation describing language (and so, I assume, obligation) seems to be hooked up with preference, a relation of what-is-better-than-what. But ordinary situations with ordinary normative constraints underdetermine such relations of what-is-better-than-what. Even so, there are plainly true sentences describing our obligations in those situations. My argument will be that this mismatch is troublemaking and that getting out of that trouble requires either giving up the direct link between obligation and preference or re-thinking the kind of things preferences can be.

2 The target

The problem I want to raise involves obligation describing language that involves expectation modals: these are modals that give voice to our obligations subject to what we know.¹ Some examples:

* Talks that this paper has grown up from date back a long time: since 2012. Thanks to everyone who helped me figure out what I was trying to say, in particular Josh Dever, Kai von Fintel, Chris Barker, Sam Carter, Andy Egan, Jeff King, Ruth Chang, Chris Kennedy, Scott Sturgeon, Peter Ludlow, Jeff Horty, Kit Fine, Ralph Wedgwood, Shyam Nair, Jan Dowell, Karen Lewis, and audiences at Northwestern University, LOGOS-Barcelona, USC, University of Chicago (twice, sorry!), NYU, the University of Arizona (twice, sorry!), the University of Toronto, the University of Texas at Austin, and the New York Philosophy of Language Workshop.

¹ See [von Fintel 2012](#) for a recent discussion of the classical approach (and its discontents) to these modals. Two notes before we get going. First: I will focus on *ought* even though there is room to wonder whether this is how we express all-in, strong obligation. That is because *ought* and *should* are weak necessity modals and it's sensible to wonder whether all-in obligation is weak in that way. (See [von Fintel & Iatridou 2008](#) and the references therein for more on weak necessity modals.) So, while I will stick with *ought*, the problem

- (1) a. I ought do the dishes.
(given the house rules plus we know that it is Friday)
- b. Jimbo ought to be in class.
(given the terms of his scholarship plus that we know school is in session)
- c. Lisa ought to go to the rally.
(given her promises plus we know she made them)
- d. Homer ought to move his car.
(given the laws plus what we know about where he parked)

A predicament is just an ordinary situation in which we know some things and face some normative constraints. The relevant reading of *ought* at stake here is sometimes called its “constrained maximization” reading: we ought to do what’s true in the best of the possibly suboptimal possibilities compatible with what we know. We can’t (that’s an expectation modal there) do what we know can’t be.²

That gives us a general template (we’ll use \square for the target modal):

Template. An obligation claim *ought b* ($\square b$) is true in a predicament iff all the best worlds given the normative constraints plus what we know in that predicament are *b*-worlds.

There are two moving pieces here: what we know and what is best. But there is no best full-stop. There are only bests relative to a given relation of better than. So *ought b* in a given predicament says that the best worlds, with re-

about preference is about whatever we use to express strong, all-in obligation. So substitute the strong expectation modal of your dialect as necessary. Second: the problem stands for deontic *oughts* of whatever flavor so long as it is tied to a relation of comparative betterness of that same flavor. I try to stay away from examples with full-on moral *oughts* so that substantive debates about moral betterness don’t distract from the structural point.

² For a contrasting reading consider:

- (i) a. Jimbo ought to be in class, but isn’t.
- b. Jimbo isn’t in class, but he ought to be.

These describe obligations that Jimbo has irrespective of what we know about his ability to, going forward, meet them. That’s why they tolerate, but the relevant readings of (1) don’t tolerate, the negation of their prejacent. These are perfectly good *oughts*, but not our topic.

spect to the relevant relation of what-is-better-than-what in that predicament and given what we know in it, are b worlds.

Our job is to fill this in by saying how the normative constraints in a predicament give us a relation of what-is-better-than-what (some kind of preference relation) and saying how that relation combines with what we know to deliver the set of best worlds. There may be no good way to do that.

3 Preference

The basic observation we will see is this: predicaments can underdetermine relations of what-is-better-than-what but nevertheless there are determinate facts about whether an *ought* is true in such situations. Coping with this mismatch is the problem. Before getting to all that, I want to have a way of framing things that is as theory-neutral as possible.

Definition 1 (Preliminaries). Fix a finite set \mathbf{A} of atomic sentences.

- i. \mathbf{L}_A^o is the smallest set containing \mathbf{A} that is closed under \neg and \cap .
- ii. \mathbf{L}_A is the smallest set including \mathbf{L}_A^o and is such that: if $a \in \mathbf{L}_A^o$ then $\Box a \in \mathbf{L}_A$; and if $a \in \mathbf{L}_A$ then $\neg a \in \mathbf{L}_A$.
- iii. $W = 2^{\mathbf{A}}$ is the set of possible worlds.

For readability I will be willfully sloppy and let a, b, c, \dots range both over non-modal (i.e., descriptive) sentences and the sets of worlds where the sentences in question are true.³ Along these lines: I will use \neg as both a connective (negation) and its set-theoretic interpretation (complementation) and use \cap both as a set-theoretic operation (intersection) and a connective that expresses it (conjunction).⁴

Take as basic the concept of a local preference: you have a local preference for b given a iff within the a -region of logical space, b is ceteris paribus better than $\neg b$. This is an expression of prima facie betterness within the

³ More carefully: (i) $\llbracket p \rrbracket = \{w : w(p) = 1\}$ for every $p \in \mathbf{A}$; (ii) for any a , $\llbracket \neg a \rrbracket = W \setminus \llbracket a \rrbracket$; and (iii) for any a and b , $\llbracket a \cap b \rrbracket = \llbracket a \rrbracket \cap \llbracket b \rrbracket$. Finally, when context allows (i.e., always) drop the scott brackets from $\llbracket a \rrbracket$ and write a .

⁴ Similarly for \cup : it represents both disjunction and its set-theoretic interpretation. Since the set of atoms is finite, so is the space of worlds; nothing much turns on this and it simplifies a few definitions and proofs. When it doesn't make much difference, let's also suppress mention of \mathbf{A} and \mathbf{L}_A even though, officially, everything is parametric on a choice of underlying language.

a -worlds. Write it this way: $b \parallel a$. So $b \parallel a$ is a claim in the metalanguage saying that there is (all else being equal) an a -preference for b over $\neg b$.

An example: an editor asks you to review a paper and you promise to do it. Suppose this is the only relevant normative constraint in the predicament. You have a local preference within the you-promised worlds: all else equal, doing the review is better than not doing it. If all else is not equal, then things may be different. That's OK.⁵

Predicaments will be represented by states.

Definition 2 (States). A state $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ is pair of a finite set \mathbf{k} of propositions and a finite set \mathbf{p} of local preferences.

The set \mathbf{k} represents what is known in \mathbf{s} and \mathbf{p} encodes the normative constraints.

Definition 3. A set \mathbf{k} of propositions is consistent iff $\cap \mathbf{k} \neq \emptyset$.

There are some minimal constraints on the space of consistent states (and thereby the predicaments they model).

Postulate 1. A state \mathbf{s} is consistent ($\mathbf{s} \neq \perp$) iff both \mathbf{k} and \mathbf{p} are consistent.

So far this only gives a necessary condition for consistency: we know what it takes for \mathbf{k} to be consistent but so far haven't said what it takes for \mathbf{p} to be consistent. We will come back to this (in Section 6).

Definition 4 (Triviality). A set S of states is trivial iff for every consistent $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ in S and distinct a, b, c at least one of the following holds:

- i. if $b \parallel a \in \mathbf{p}$ then $b \parallel a \cap c \notin \mathbf{p}$;
- ii. if $b \parallel a \in \mathbf{p}$ then $c \parallel a \notin \mathbf{p}$;
- iii. if $b \parallel a \in \mathbf{p}$ then $\neg b \parallel c \notin \mathbf{p}$;

So if the space of states is non-trivial then local preferences can in principle overlap in both arguments and can conflict and be over-ridden. I will assume that normative constraints are like that and so the thing we are here taking them to express (local preferences) are, too.

⁵ You can read [Horty 2012](#) (especially Chapters 3 and 4) as developing a theory of what you ought to do in predicaments. Notably, the framework there is not based on the template that what you ought to do is what is best given what you know with respect to a relation of what is better than what.

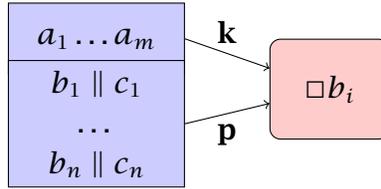


Figure 1 Predicament: \mathbf{k} (set of propositions) + \mathbf{p} (set of local preferences)

Postulate 2. The set \mathbf{S} of states is non-trivial.

Definition 5 (Support). Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ be any state. For any a and b :

- i. $\mathbf{s} \models a$ iff $\cap \mathbf{k} \subseteq a$;
- ii. $\mathbf{s} \models \neg \square b$ iff $\mathbf{s} \not\models \square b$.

Support will play the role of truth in what follows. Descriptive sentences are not bivalent: \mathbf{s} can fail to support a and fail to support $\neg a$. This makes sense since you and hence your states are in general only partially informed about the facts. Not so for obligation: given your information and local preferences, either you ought to do a or not. This definition doesn't yet say what it takes for sentences like $\square b$ to be supported or true in a state but it does put constraints on possible analyses. The analyses we will consider below are ways of filling in the missing clause in this definition; with them in place we will have candidate support relations.⁶

Now to the problem. Take the ordinary predicament above: you know that you promised to do the review (a) and the the only normative constraint is that it is ceteris paribus better to do it given your promise ($b \parallel a$). In a state like this you have chunky preferences: you have definite attitudes within the a worlds about the desirability of b versus $\neg b$ but you don't have any attitudes about an unrelated c given a and you don't have attitudes about b simpliciter. The normative constraints are unspecific in these ways.⁷ Still, there are *oughts* that are clearly true.

⁶ When \models occurs without subscripts, it is either unspecific (what is said goes or ought to go for any supports relation) or is whichever candidate relation is under discussion; context should disambiguate.

⁷ There is an analog with credence: what to do if we have imprecise or unspecific information about whether a (*The chance of rain is between 50-70%*)? Our credences (some say) should be similarly mushy (see, for instance, [Joyce 2011](#)). As we'll see the standard way to understand

- (2) (All else equal) Better to review (than not), given you promised. $b \parallel a$
 You promised. a
 You ought do the review. $s \models \Box b$

Having chunky preferences is a kind of indeterminacy. But the way those normative constraints pull on us apparently is not thereby indeterminate.

We will look at three ways of dealing with the mismatch between indeterminacy of predicaments and the determinacy of what we ought to do in them. The would-be solutions seamlessly fit into the template by invoking a more familiar sort of thing: a global preference ordering. This is no accident. Such relations are well-behaved, well-understood, and seem to be what's required to determine what's best.⁸

Definition 6 (Global preference ordering). A global preference ordering \preceq is a transitive relation over the set W of worlds. It is a weak global preference ordering if it is also reflexive. The strict part of \preceq is $<$ where $w < v$ iff $w \preceq v$ but $v \not\preceq w$.

The best worlds in a set x with respect to \preceq are those in x that are not bettered by any others in x .

Definition 7 (Best). Fix a global preference ordering \preceq .

$$\text{best}_{\preceq}(x) = \{w \in x : \text{there is no } v \in x \text{ such that } v < w\}$$

Every non-empty x has some non-bettered worlds in it since every finite set is well-ordered by the strict part of \preceq .

Such global orderings (seem to) contain a lot more information about comparative goodness than your run of the mill set of local preferences do. This is true even if the ordering contains incomparable worlds: even if x and y aren't comparable to each other they will typically still be comparable with to bunches of other possibilities and occupy a determinate spot in the ordering. Still, since given such an ordering it is straightforward to find the best worlds in a set with respect to it, the hope is that global preference orders can be leveraged to represent the normative information in a set of

mushy credences won't work for chunky preferences. That is somewhat surprising since generally what goes for credence goes for sufficiently rich preferences and vice versa.

⁸ Such orderings (for now) need not be connected: there may be a w and v such that $w \not\preceq v$ and $v \not\preceq w$. Later we will mostly be concerned with sets of preference orderings that are both total (no incomparabilities) and strict (no ties).

local preferences, bridging the gap between local preferences and the *oughts* that are true based on them.

4 Preference determination

A sets of local preferences is chunky and unspecific and indeterminate in ways that a global preference ordering isn't. Perhaps the way to deal with this indeterminacy is simple: perhaps every set of local preferences, indeterminate though it may be, nevertheless determines a global ordering in a reasonable way. The good news here is there is a natural and principled and well-traveled route to take.

Definition 8 (Flouting, complying, confirming). Let $b \parallel a$ be any local preference. A world w flouts $b \parallel a$ iff $w \in (a \cap \neg b)$. A world w complies with $b \parallel a$ iff w doesn't flout $b \parallel a$. A world w confirms $b \parallel a$ iff $w \in a$ and w complies with $b \parallel a$.

The set $a \cap \neg b$ is the flouting proposition for $b \parallel a$ and $\neg a \cup b$ is its complying proposition. A privileged subset of the complying proposition for $b \parallel a$ is made up of those worlds that are also in a : these constitute $a \cap b$ and make up the confirming proposition for $b \parallel a$.

Definition 9 (Induced preference ordering). Let \mathbf{p} be a (finite) set of local preferences. The global preference ordering induced by \mathbf{p} is the ordering $\preceq_{\mathbf{p}}$ such that for any w, v :

$$w \preceq_{\mathbf{p}} v \text{ iff } \{b \parallel a \in \mathbf{p} : v \in (\neg a \cup b)\} \subseteq \{b \parallel a \in \mathbf{p} : w \in (\neg a \cup b)\}$$

That is: $w \preceq_{\mathbf{p}} v$ iff w complies with every local preference in \mathbf{p} that v does. The idea is not new.⁹ What we have done is to treat local preferences as the ingredients of premise sets or, as they are now known, ordering sources for deontic modals.

This naturally pairs with the template.

⁹ See Lewis 1981 where this way of inducing orderings is used to show that the difference between premise semantics (Veltman 1976, Kratzer 1981b) and ordering semantics (Stalnaker 1968, Stalnaker & Thomason 1970, Lewis 1973, Pollock 1976) for counterfactuals is exaggerated: as far as evaluating counterfactuals is concerned, a premise set and its induced ordering come to the same thing.

Analysis 1. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ be any state.

$$\mathbf{s} \models_{\perp} \Box b \text{ iff } \text{best}_{\preceq_{\mathbf{p}}}(\cap \mathbf{k}) \subseteq b$$

where $\preceq_{\mathbf{p}}$ is the weak global preference induced by \mathbf{p} .

This is a multi-step procedure: derive orderings from the set of complying propositions for \mathbf{p} (an ordering source); find the best worlds in that ordering, modulo \mathbf{k} (a modal base); *ought* universally quantifies over this set.¹⁰

This is a candidate solution to our problem: it says how to go from underdetermined predicaments to determinate truth conditions for *oughts*. It has a lot going for it. It rightly predicts that *oughts* can come and go both depending on what you know and depending on what the normative constraints are. For instance: suppose you are a committed promise-keeper but that you haven't made a promise to go to the pub.¹¹ Then you are unconstrained: you can do whatever, pubwise.

- | | |
|---|--------------------------------------|
| (3) Better to go to the pub given you promised. | $b \parallel a$ |
| It's not the case that you ought go to the pub. | $\mathbf{s} \not\models \Box b$ |
| It's not the case that you ought not go to the pub. | $\mathbf{s} \not\models \Box \neg b$ |

Things are different when you (know you) made a promise, though.

- | | |
|---|-----------------------------|
| (4) Better to go to the pub given you promised. | $b \parallel a$ |
| You promised to go to the pub. | a |
| You ought go to the pub. | $\mathbf{s} \models \Box b$ |

Analysis 1 predicts this pattern smoothly. Here's why (this is all pictured in Figure 2(a)).¹² In the case of (3) the lone local preference induces an

¹⁰ This implements in the current framework the standard-bearer in semantics for all sorts of relative modality: see Kratzer 1981a, 1991, 2012.

¹¹ I sometimes elide the "all else equal" qualifier on local preferences from here on out; it's always there in spirit. Here and in some of the following examples (some of) the local preferences are unconditional: that is, conditional on \top . This is fine: the arguments have to do with the structure of local preferences and their allegedly induced global orderings and that is somewhat easier to track when there are only a few propositions in play. If you find things (even) more natural for constraints that are genuinely conditional, be my guest in crafting examples with three basic propositions in play instead.

¹² Conventions for the graphs: \preceq is the reflexive transitive closure of the depicted arrow-relation (thus $w \preceq v$ iff there is a directed path of length 0 or more from w to v); $w \notin \cap \mathbf{k}$ iff w 's node is grayed out.

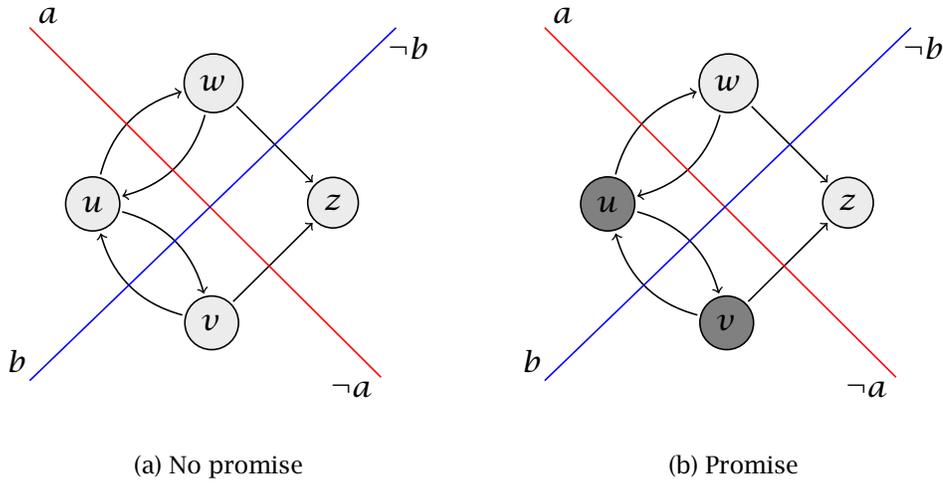


Figure 2 Promise keeping

ordering \preceq_p that divides the worlds into two clumps: those that comply with the preference (those in $\neg a \cup b$) and those that flout it (those in $a \cap \neg b$). Each clump is an equivalence class with respect to \preceq_p and each member of the first clump is strictly better than any in the second. Since you don't know anything relevant ($\mathbf{k} = \emptyset$) it follows that $\text{best}_{\preceq_p}(\cap \mathbf{k})$ contains worlds in $a \cap b$ and worlds in $\neg a \cap \neg b$ and so \mathbf{s} neither forces $\Box b$ nor $\Box \neg b$ and so instead supports their negations.

Things are different in (4). Now you know something (a) that interacts with your constraints. Now, pick any world in a that is not bettered by some other world in a . (In Figure 2(b): this is $\{w\}$.) Is it a b -world? Yes, it is. Knowing even more could, in principle, remove the obligation: if the appointed hour comes and goes and you don't go to the pub, then you will know $\neg b$ and the best worlds compatible with this will not be b worlds.

The general feature this is an instance of: what's supported isn't persistent in what you know.

Definition 10 (Persistence). For any $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ and $\mathbf{s}' = \langle \mathbf{k}', \mathbf{p}' \rangle$: \mathbf{s}' extends \mathbf{s} ($\mathbf{s} \leq \mathbf{s}'$) iff $\mathbf{k} \subseteq \mathbf{k}'$ and $\mathbf{p} \subseteq \mathbf{p}'$.

- i. A sentence a is persistent in \mathbf{k} (alternatively: in \mathbf{p}) with respect to \models iff $\mathbf{s} \models a$ and $\mathbf{s} \leq \mathbf{s}'$ imply $\mathbf{s}' \models a$ where $\mathbf{p} = \mathbf{p}'$ (alternatively: where $\mathbf{k} = \mathbf{k}'$).

- ii. A sentence a is persistent (full-stop) with respect to \models iff it is persistent in both \mathbf{k} and \mathbf{p} .

Non-persistence is a kind of nonmonotonicity and both $\Box b$ and $\neg\Box b$ exhibit it: in particular, knowing more can both usher in and sweep out obligations.¹³

Proposition 1. Both $\Box b$ and $\neg\Box b$ are non-persistent in \mathbf{k} with respect to \models .

Obligations are likewise non-persistent in local preferences. An example: suppose the only normative constraint is that you (unconditionally) prefer becoming mayor to not and the only relevant information is that going to the pub is a necessary condition for becoming mayor. Well, then, you ought to go to the pub.

- | | |
|---|-----------------------------|
| (5) Better to become mayor than not. | $a \parallel \top$ |
| Either you don't become mayor or you go to the pub. | $\neg a \cup b$ |
| You ought to go to the pub. | $\mathbf{s} \models \Box b$ |

Suppose instead that you face the additional constraint: you (unconditionally) prefer to not go to the pub.¹⁴

- | | |
|--|---------------------------------------|
| (6) Better to become mayor than not. | $a \parallel \top$ |
| Better to not go to the pub than to go. | $\neg b \parallel \top$ |
| Either you don't become mayor or you go to the pub. | $\neg a \cup b$ |
| It's not the case that you ought to go to the pub. | $\mathbf{s}' \not\models \Box b$ |
| It's not the case that you ought to not go to the pub. | $\mathbf{s}' \not\models \Box \neg b$ |

This is too bad for you: your local preferences compete and pull you in opposite directions. As a result you don't have the obligation to go to the pub and you don't have the obligation to not go to the pub. Though you ought to do one or the other.¹⁵

¹³ The main text includes (hopefully) enough detail to see why the results reported hold good. For those who want more, see the appendix.

¹⁴ The example (based on one in [Kratzer 1981a](#)) is an instance of the "Nixon diamond" in nonmonotonic logic ([Reiter & Criscuolo 1981](#)). Suppose your information is that Quakers are (normally) pacifists and that republicans are (normally) not pacifists. If all you know about Nixon is that he is a Republican and a Quaker, then what should you conclude about his being a pacifist?

¹⁵ Thus when it comes to (potential) moral conflicts, the account is an implementation of what [Horty \(2012\)](#) calls the "disjunctive account". See also [Gillies 2014](#), [Horty 2014](#).

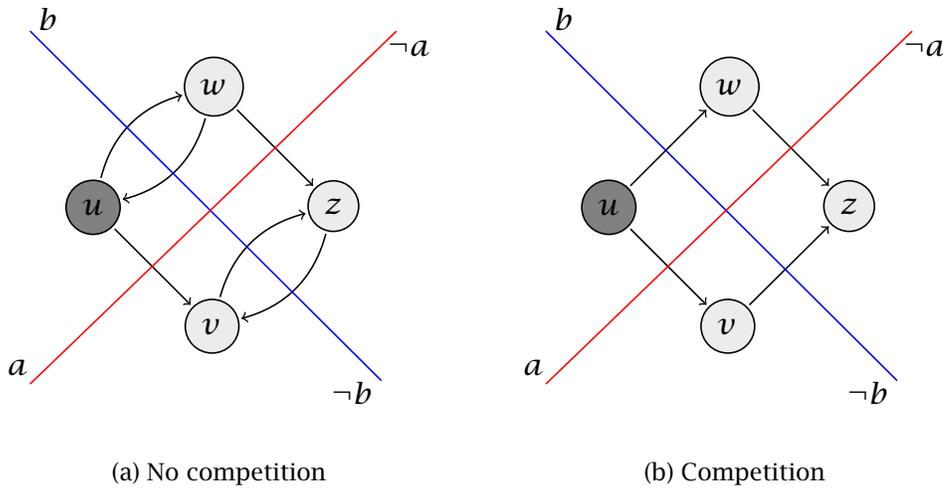


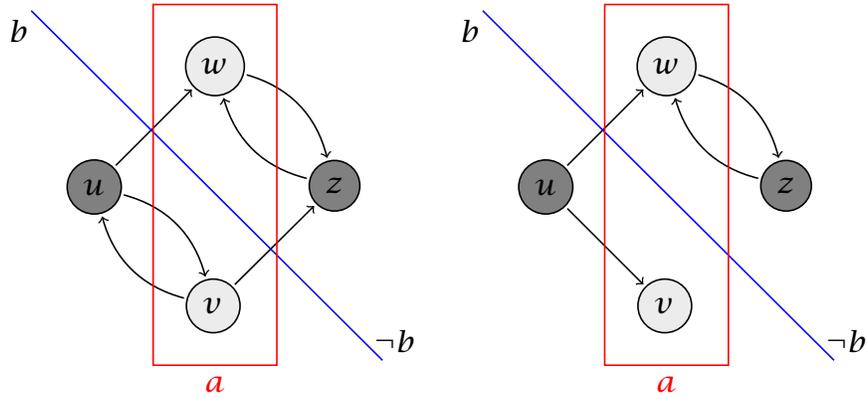
Figure 3 Becoming mayor

Again, Analysis 1 predicts this pattern smoothly. When the only local preference is $a \parallel \top$ then as far as \leq_p is concerned there are the complying worlds and the flouting worlds, where the compliers are each as good as each other and strictly better than each flouter and each flouter is equally as good (or bad, as it happens) as each other. In the ordering in Figure 3(a), u is the only world ruled out by \mathbf{k} . There is a path from w to v but not vice versa, so $\text{best}_p(\cap \mathbf{k}) = \{w\}$ and so $\mathbf{s} \Vdash_{\top} \Box b$.

When you instead have the local preferences $\mathbf{p}' = \{a \parallel \top, \neg b \parallel \top\}$, worlds that were previously tied with respect to \leq_p no longer are (the ordering in Figure 3(b)): u complies with both local preferences, w and v comply with one each, and z flouts both. Now w and v are incomparable: there is no path from w to v and no path from v to w . But your information hasn't changed. Becoming mayor still requires, as a matter of brute fact, that you go to the pub. That rules out u (and nothing else). Hence the best worlds compatible with \mathbf{k} according to \leq_p are w and v . Since one is a b -world and one a $\neg b$ -world it follows that $\mathbf{s}' \not\Vdash_{\top} \Box b$ and $\mathbf{s}' \not\Vdash_{\top} \Box \neg b$. That gives us the following result.

Proposition 2. Both $\Box b$ and $\neg \Box b$ are non-persistent in \mathbf{p} with respect to \Vdash_{\top} .

Knowing more can add to and reduce your obligations and so can becoming acquainted with new normative constraints. Analysis 1 gets this right.



(a) No conflict

(b) Unwanted conflict

Figure 4 Specificity

5 Preference (un)faithfulness

As tidy as this all seems, I don't think taking local preferences to determine global preference orderings in this way is right and therefore a candidate solution like Analysis 1 is not right.

Local preferences express prima facie betterness. So often enough trade-offs have to be made in finding what is best. The mechanism that induces the orderings adjudicates trade-offs between local preferences in a way that over-generates incomparabilities between worlds. It lumps predicaments in which one local preference over-rides another together with predicaments where there is a stand-off between competing local preferences. In such predicaments it thus consults orderings that, intuitively, it shouldn't.

Predictably, this has empirical consequences. Starting there: (7) and (8) form a minimal pair.

(7)	(All else equal) Better to go to the pub than not.	$b \parallel \top$
	It is Sunday.	a
	You ought to go to the pub.	$s \models \square b$

(8)	(All else equal) Better to go to the pub than not.	$b \parallel \top$
	(All else equal) Better to not go to the pub given it's Sunday.	$\neg b \parallel a$

It is Sunday.

You ought not go to the pub.

a
 $s \models \square \neg b$

Figure 4(a) is a picture of the predicament in (7). There is no path from w to v but there is a path from v to w . Thus the best world in a is v . This is a pub-going world, so you ought to go to the pub. This is right.

In the predicament in (8), you ought not go to the pub but Analysis 1 does not agree. It agrees that it is not the case that you ought to go, but it does not say that you ought to not go.¹⁶

You can see this in Figure 4(b): the induced preference relation now treats w and v as incomparable. Each complies with one local preference and flouts the other. Since $a = \{w, v\}$, both are in $\text{best}_{\leq_p}(\cap \mathbf{k})$. But $w \in \neg b$ and $v \in b$ and so both $s \not\models \square b$ and $s \not\models \square \neg b$. There is a predicted stand-off between your local preferences and so an under-generation of obligations.¹⁷

¹⁶ It matters that there is nothing more to the two local preferences. For instance, if instead each local preference comes as the command of one of two equally in-charge deities, maybe things are different:

- (i) God A commands: better to go to the pub than not.
God B commands: better not to go to the pub given it's Sunday.
It is Sunday
It's not the case you ought to go to the pub.
It's not the case you ought to stay away.

Here, it seems, there is conflict. But the source of the local preferences is where the conflict lies and so that needs representing. (The commands of deities aren't all that *ceteris paribus*.) Something like this: $b \parallel b_A$ and $\neg b \parallel \neg b_B$ where the triggering propositions are which god commanded which thing. However this gets done, neither trigger entails the other and so it is not relevantly like example (8).

¹⁷ Tempting thought: this problem for Analysis 1 is a (generalized) version of the problem that the Miners' puzzle adapted to this new setting (Kolodny & MacFarlane 2010). There the problem is that where a, b, c are all compatible it is hard to predict that *if* a , *ought* b and *if* $a \cap c$, *ought* $\neg b$ can be true together. There's some resemblance here but don't lean too much on it. As we'll see shortly, there's no tension between $b \parallel a$ and $\neg b \parallel a \cap c$ in the precise sense that there is an ordering that is faithful to them both. The problem with Analysis 1 isn't a dearth of orderings but that it consults the wrong ones. In a bit (footnote 20) we'll see that amending Analysis 1 by allowing the derived orderings to be information dependent still doesn't get things right. That said: the pattern in (8) and Proposition 3 does have an exact parallel elsewhere. It is called the "Penguin principle" in nonmonotonic logic: given only the information that Tweety is a bird you conclude (defeasibly) Tweety flies, but given the additional information that Tweety is a penguin you take that back and conclude that, no, Tweety does not fly (Etherington & Reiter 1983).

The official and slightly more general form of the problem:

Proposition 3. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a, c\}$ and $\mathbf{p} = \{b \parallel a, \neg b \parallel a \cap c\}$. Then $\mathbf{s} \not\vdash_1 \Box \neg b$.

To be clear: it is not just that this analysis gets the wrong verdict here. (Though, come on, it does.) Maybe, in fact, you disagree on the judgment in (8) that you ought to stay away from the pub. (Though, come on, you don't.) Set that aside. Analysis 1 takes the stance that there are no genuine predicaments in which one local preference over-rides another. They are all, in the end, stand-offs between competing local preferences that result in incomparability. That is a bold stance.

Over-riding seems to require one local preference to have a triggering condition that is more specific than another's. The local preference to stay away from the pub on Sundays should over-ride the local preference for pub going. So maybe we can patch Analysis 1: anytime \mathbf{p} contains a local preference like $b \parallel a$ and one like $\neg b \parallel a \cap c$ with a strictly stronger triggering condition, the induced ordering should only pay attention to the more specific $\neg b \parallel a \cap c$. So, in (8), since a asymmetrically entails \top the induced ordering should only pay attention to $\neg b \parallel a$ and it should ignore $b \parallel \top$.

This is not the way to go, for two reasons. It is a special kind of ad hoc. It requires that we intervene on behalf of the analysis, guiding it to pay attention to all and only the right local preferences in constructing an ordering to represent the normative information in a predicament. We, not the mechanism for inducing the ordering, are doing all the heavy lifting of adjudicating trade-offs between local preferences.¹⁸

And it doesn't get at the root issue anyway. Over-riding is a matter of taking precedence. The induced preference mechanism wrongly treats precedence-taking, whether or not the shape of the predicament is exactly like the Penguin principle, as symmetric conflict. You can see this in predicaments where the precedence-taking is a matter of how what you know interacts with your total set of local preferences.

¹⁸ There are parallel patches in some early nonmonotonic logics that also fail to smoothly predict the penguin principle (e.g., constrain the order in which defaults are applied! prioritize which predicates are minimized! and so on) and Asher & Morreau (1991: 388) are not having any of it, saying that such theories with such patches "commit themselves to the Hypothesis of the Ghost in the Machine" of default reasoning. Instead, they say, the penguin principle should "emerge naturally". Hard agree.

An example: you have a favorite aunt. Sometimes you visit her and sometimes you call her. It is all else equal better to call before you go for a visit ($b \parallel a$). You are going for a visit (a). Well, then, you ought to call ($\Box b$). But a surprise party changes things: it's better to visit her given that there is a surprise party but in that case it's better to not call first.

- (g) Better to call before you go, given that you're going to visit. $b \parallel a$
 Better to not call but visit given that there's a surprise party. $\neg b \cap a \parallel c$
 You are going to visit and there's a surprise party. $a \cap c$
 You ought to not call before. $\mathbf{s} \models \Box \neg b$

Three observations. First: $\neg b \cap a \parallel c$ seems to over-ride $b \parallel a$ here in just the same way that $\neg b \parallel a$ over-rides $b \parallel \top$ in (8). In fact the Penguin principle on display in (8) is what you get from the pattern in (9) when $a = \top$. So, just like you ought to stay away from the pub, you ought not call. Second: neither local preference has a triggering condition that, alone or together with what you know, asymmetrically entails the other. Third: Analysis 1 does not say you ought to not call and for exactly the reason we've already seen. It wrongly treats the relationship between $b \parallel a$ and $\neg b \cap a \parallel c$, when you know $a \cap c$, as a case of conflict leading to a stand-off in the set of best worlds compatible with what you know.¹⁹

Proposition 4. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a, c\}$ and $\mathbf{p} = \{b \parallel a, \neg b \cap a \parallel c\}$. Then $\mathbf{s} \not\models_1 \Box \neg b$

Consider a minimal model where $\mathbf{A} = \{a, b, c\}$. No world can confirm both $b \parallel a$ and $\neg b \cap a \parallel c$. So any world compatible with $a \cap c$ that complies with one of the local preferences has to flout the other. Figure 5 shows the situation: in Figure 5(a) shows what you get from $b \parallel a$ on its own (as a point fo reference) and Figure 5(b) has the ordering we care about that is induced by $\mathbf{p} = \{b \parallel a, \neg b \cap a \parallel c\}$. There are three (color coded!) equivalence classes of worlds. The worlds in $[u] = \{s, u, z\}$ comply with both $b \parallel a$ and $\neg b \cap a \parallel c$. The worlds in $[w] = \{s, w, x\}$ comply with just $b \parallel a$ and the worlds in $[v] = \{t, v\}$ comply with just $\neg b \cap a \parallel c$. Every world in $[u]$ is strictly better than every world in $[w]$ and strictly better than every world in $[v]$. But no world in $[w]$ is comparable to any world in $[v]$ and so in

¹⁹ Bonus observation: the Analysis 1 + patch package still declares that predicaments with over-riding preferences are impossible. Still seems bold to rule these out a priori.

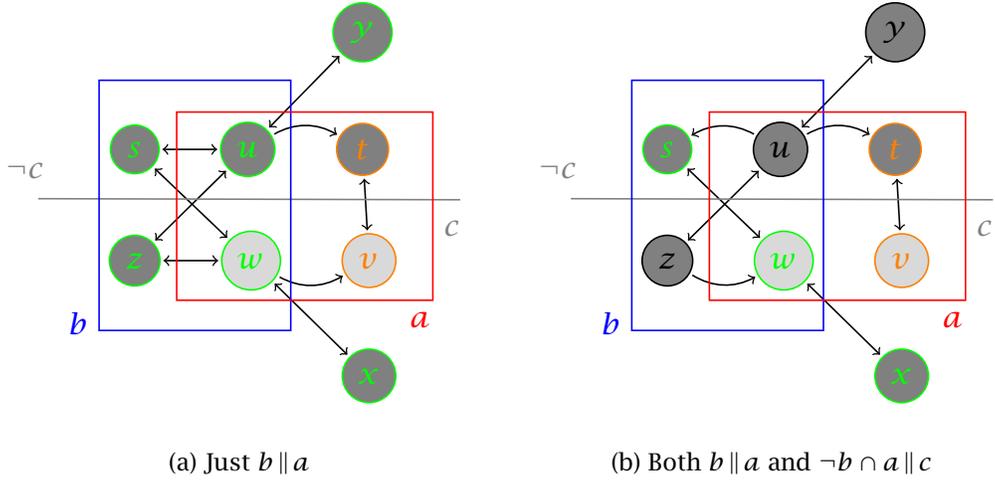


Figure 5 Over-generating conflict in $\mathbf{k} = \{a \cap c\}$

particular $\text{best}_{\leq_p}(\cap \mathbf{k}) = \{w, v\}$. Thus, according to Analysis 1, you do not have the obligation to not call.²⁰

Lumping predicaments with over-riding constraints together with predicaments with conflicting constraints is wrong, but the reason for the lumping is more wrong. One way of exposing the trouble focuses on the linearizations of the induced (weak) global preference ordering.

Definition 11 (Strict total global preference). A strict total preference ordering $<$ is a global preference ordering that is also trichotomous.²¹

These differ in two important ways from our earlier global relations: these are strict (no ties) and are total (no incomparabilities).²²

²⁰ The problem is not unique to Analysis 1. The problem lies with the Lewisian mechanism for inducing global orderings from sets of propositions and therefore is inherited by refinements of that method. For instance, the mechanism in Cariani et al. 2013 is meant to deliver different global orderings depending on what information is present in a predicament by partitioning the worlds in a way based on it. It does, but under the natural partition of actions in (8) it, too, delivers incomparability between the best (given what you know) pub-going worlds and pub-avoiding worlds.

²¹ Trichotomy: for any w, v exactly of one of the following holds: $w < v$, $w = v$, or $v < w$. This implies comparability.

²² Given such a $<$, define its weak counterpart as follows: for any w and v , $w \leq v$ iff $w < v$ or $w = v$. The resulting relation is a weak total ordering of W . We will no longer have any

Definition 12 (Linearizations). Let \preceq be any global preference ordering. A strict total preference ordering \prec' linearizes \preceq iff: if $w \prec v$ then $w \prec' v$. If \prec' linearizes \preceq it is a linearization of \preceq .

Intuitively: linearizations of \preceq agree with the strict part of it and break all ties and resolve all incomparabilities.²³ But not all ways of breaking ties and resolving incomparabilities in \preceq will result in a proper strict ordering that obeys trichotomy: some ways violate transitivity, some introduce new incomparabilities, and some don't agree with the strict part of \preceq . So you might wonder if every weak global preference ordering has a linearization. Yes: by successively picking a best world according to \preceq in a sequence of ever so slightly smaller sets of worlds, you can create a linearization of \preceq . The proof is constructive. It exploits the fact that we can construct ranking functions that respect the strict part of \preceq .

Definition 13 (Ranking functions). A ranking function is a function $\kappa: W \rightarrow \mathbb{N}_{\geq 0}$ such that $\kappa^{-1}(0) \neq \emptyset$.

Think of $\kappa(w)$ as the qualitative degree of suboptimality of w .²⁴ Construct κ as follows. Pick a w_1 that is not bettered (with respect to \prec) by any other world in W ; assign $\kappa(w_1) = 0$. Next pick any w_2 in $W \setminus \{w_1\}$ that is not bettered by any world in $W \setminus \{w_1\}$; assign $\kappa(w_2) = 1$. Lather, rinse, repeat until you finally assign $\kappa(w_n) = n - 1$. Finally, define $w \prec v$ iff $\kappa(w) < \kappa(v)$. This strict total ordering agrees with the strict part of \preceq , breaks its ties, and resolves its incomparabilities. That is:

Theorem 1. Every weak global preference ordering has a linearization.

use for strict preferences that are not also trichotomous so we will drop the “total” qualifier unless context demands it.

²³ Linearizations, like the mechanism for inducing preferences, also have conditional roots: in [van Fraassen 1974](#) and [Lewis 1981](#) they are called “Stalnaker refinements” and play a role in showing that there isn't much disagreement between versions of ordering semantics that allow ties and incomparabilities (like [Pollock's](#)) and versions that allow ties but no incomparabilities (like [Lewis's](#)) nor between the yes-to-ties-no-to-incomparabilities versions and those that allow neither (like [Stalnaker & Thomason's](#)). Linearizations also play a role least-commitment planning (see, for instance, [Pollock 1998](#)).

²⁴ Ranking functions are a way of ordering worlds and propositions by associating them with non-negative integers: the lower the rank, the better. They are developed as a qualitative model of belief and belief dynamics in [Spohn 1988](#). For a different use of them in the deontic setting see [Spohn forthcoming](#).

In general an induced $\preceq_{\mathbf{p}}$ has a bunch of linearizations that differ in how they settle each tie or incomparability in $\preceq_{\mathbf{p}}$.²⁵ Linearizations seem way more determinate than the weak (and possibly unconnected) preferences that they linearize and so may seem out of place here. We will see that that is not so.

The way of determining the best worlds in a set x is agnostic on whether bestness is based on a weak global preference relation or a strict and total global preference relation. Doing this across linearizations shows that it doesn't matter whether we talk about what's best given a weak and possibly unconnected global preference relation or about what's best in all of its linearizations.

Proposition 5. Fix a set \mathbf{p} of local preferences and its induced global ordering $\preceq_{\mathbf{p}}$. Let \prec_1, \dots, \prec_n be the linearizations of $\preceq_{\mathbf{p}}$. Then for any set x :

$$\text{best}_{\preceq_{\mathbf{p}}}(x) = \bigcup_{i=1}^n \text{best}_{\prec_i}(x)$$

We therefore lose nothing by talking about a set of linearizations of $\preceq_{\mathbf{p}}$ rather than the induced $\preceq_{\mathbf{p}}$ itself. Analysis 1 can't tell the difference.

So return to (8): we know that $\mathbf{s} \not\perp_{\mathbf{1}} \square \neg b$. In terms of the linearizations of $\preceq_{\mathbf{p}}$, this is because there is a linearizing \prec according which the best worlds in a are in $\neg b$ but also a linearizing \prec' according which the best worlds in a are in b (Figure 5). But wait: the whole idea is that you have a local preference $\neg b \parallel a$: within the a -region, $\neg b$ is supposed to be ceteris paribus better. The preference ordering \prec' just doesn't get you. The same goes for (9). In each case, there is a linearization that doesn't merely go beyond the normative information in \mathbf{p} ; in a precise sense it is unfaithful to it.

To say just what that amounts to, first lift the strict global ranking \prec on worlds to one on propositions.

Definition 14 (Lifted preferences). For any strict total global preference ordering \prec its propositional lift \preceq^+ on $\wp(W)$ is the relation:

$$a \preceq^+ b \text{ iff for any } w: \text{ if } w \in b \setminus a \text{ then } v \prec w \text{ for some } v \in a$$

The strict part of \preceq : $a \prec^+ b$ iff $a \preceq^+ b$ and $b \not\preceq^+ a$.

²⁵ This is not far off from the Szpilrajn extension theorem that every preorder has a complete extension (Szpilrajn 1930). Linearizations aren't always extensions of the orderings they linearize; instead, they extend sub-relations of those orderings.

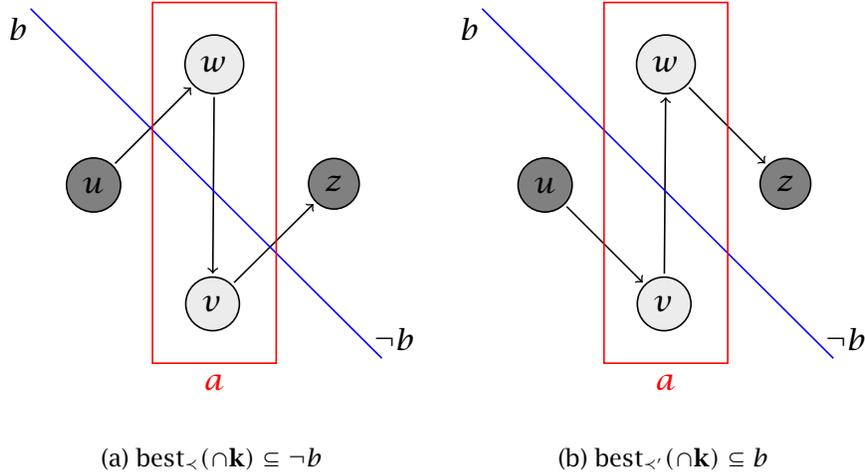


Figure 6 Two linearizations of $\preceq_{\mathbf{p}}$

So a is better than b just when every $b \cap \neg a$ world is bettered by some a world or other and not vice versa. Note that \preceq^+ is an economist-approved, weak preference ordering over propositions (it is transitive and connected).²⁶

Definition 15 (Faithfulness). A strict global preference ordering $<$ is faithful to $b \parallel a$ iff $a \cap b < a \setminus b$; it is faithful to a set \mathbf{p} of local preferences iff it is faithful to each $b \parallel a \in \mathbf{p}$.

Being faithful to $b \parallel a$ means ranking the confirming proposition $a \cap b$ as better than the flouting proposition $a \cap \neg b$. This is different from the induced ordering mechanism.

Proposition 6. A strict global preference ordering $<$ is faithful to $b \parallel a$ iff $\text{best}_{<}(a) \subseteq b$.

The issue with induced preference orderings is that they can have linearizations which are not faithful to the set of local preferences that induce them. Using that mechanism, as Analysis 1 does, obligation describing language can end up appealing to global orderings that, intuitively, it shouldn't. That is the case in predicaments where one local preference should over-ride another one.

²⁶ When it is clear from context whether $<$ or $<^+$ is in play, omit the superscript.

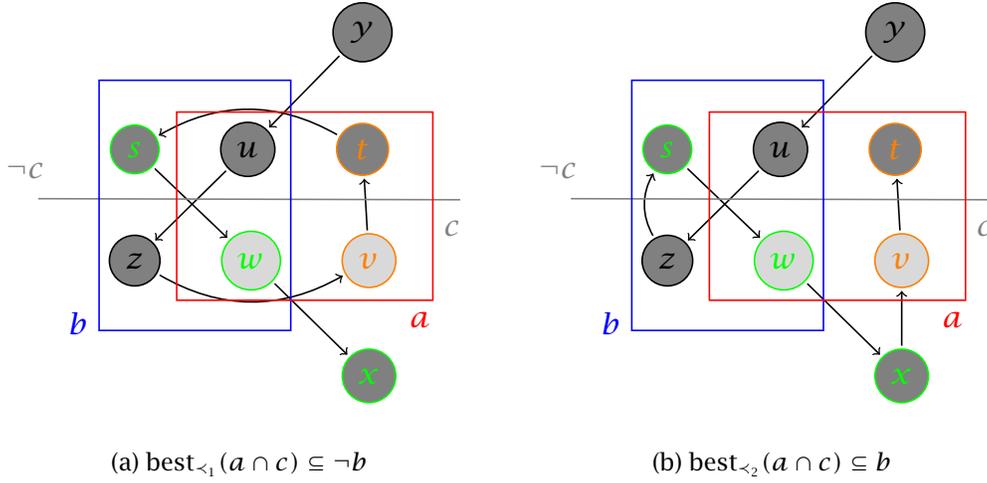


Figure 7 Two linearizations of $\preceq_{\mathbf{p}}$

So, for example, in example (8) where $\mathbf{p} = \{b \parallel \top, \neg b \parallel a\}$ there are two linearizations of \mathbf{p} : one which is faithful to $\neg b \parallel a$ and one which isn't faithful to $\neg b \parallel a$ (see Figure 5). The unfaithful one is troublemaking: according to it, the best worlds in a are not in $\neg b$. Similarly in example (9). Where $\mathbf{p}' = \{b \parallel a, \neg b \cap a \parallel c\}$, there are multiple linearization of $\preceq_{\mathbf{p}'}$: some faithful to $\neg b \cap a \parallel c$ and some unfaithful to it (Figure 7). It is the second one that is troublemaking. According to it, the best worlds in $a \cap c$ are not in b . Officially:

Proposition 7. Let $\mathbf{p} = \{b \parallel \top, \neg b \parallel a\}$. There is a linearization of $\preceq_{\mathbf{p}}$ that is unfaithful to \mathbf{p} . Similarly, where $\mathbf{p}' = \{b \parallel a, \neg b \cap a \parallel c\}$, there is a linearization of $\preceq_{\mathbf{p}'}$ that is unfaithful to \mathbf{p}' . Moreover, where $<^x = < \cap (x \times x)$, there are linearizations of $\preceq_{\mathbf{p}'}$, $<_1$ and $<_2$, such that $<_1^a$ and $<_2^c$ are unfaithful to \mathbf{p}' .

Using induced global preference orderings gets things wrong in two ways. First: the induced orderings carry more information than the thing they are modeling. In the same way that a single probability function isn't cut out to model an agent with mushy credences, a single global preference ordering isn't cut out to model an agent with chunky preferences. Trying to force them to do this job forces them to do it in a way that is unfaithful to the local preferences. Second: the induced orderings carry less information than the thing they are modeling. The induced orderings lump together

situations in which a more specific local preference trumps a less specific one with situations in which local preferences genuinely compete. The empirical upshot is undergenerating obligations in those cases.

6 Constraining preferences

Having chunky preferences is a kind of indeterminacy. Analysis 1 goes wrong by insisting that, appearances to the contrary, chunky preferences do not underdetermine proper global counterparts. The way it does this can end up treating as relevant global orderings that are unfaithful to the constraints in predicaments. This is not how to deal with indeterminacy. Perhaps, instead, we should cope with it the way we cope with other forms of indeterminacy: by quantifying over ways of making things more determinate. So what you ought to do is not what is best with respect to the relevant global preference ordering but what is best with respect to all relevant orderings. Which orderings count as relevant? The ones faithful to the local preferences.²⁷

Analysis 2. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ be any state.

$$\mathbf{s} \models_{\frac{1}{2}} \Box b \text{ iff } \text{best}_{<}(\cap \mathbf{k}) \subseteq b$$

for every strict global preference $<$ faithful to \mathbf{p} .

Analysis 1 and Analysis 2 share a commitment to representing local preferences by appeal to global surrogates. But they otherwise differ in worldview: $\models_{\frac{1}{2}}$ treats predicaments as constraining, and not in general determining, the global orderings relevant for *oughts*.

This gets a lot right. It embraces the idea that local preferences underdetermine a proper global preference ordering. But it copes with it gracefully, agreeing with a lot of the verdicts that Analysis 1 gets right without consulting wayward orderings.

Take, for starters, the predicament in (4): you have the local preference to go to the pub given that you promised ($b \parallel a$) and the information that you promised (a). Both Analysis 1 and Analysis 2 agree: you ought to go to the pub. For Analysis 2 the reason is clear. Take any $<$ faithful to \mathbf{p} : since it is

²⁷ This strategy, as we'll see, also has conditional roots (in particular, it is related to the pairing of Stalnaker-Thomason-like orderings with supervaluations (van Fraassen 1974) and (if you squint) has a connection to ways of coping of with vagueness (Fine 1975) and mushy credences (Joyce 2011). These family resemblances make a lot of sense since the idea is that predicaments are shot through with a specific kind of indeterminacy.

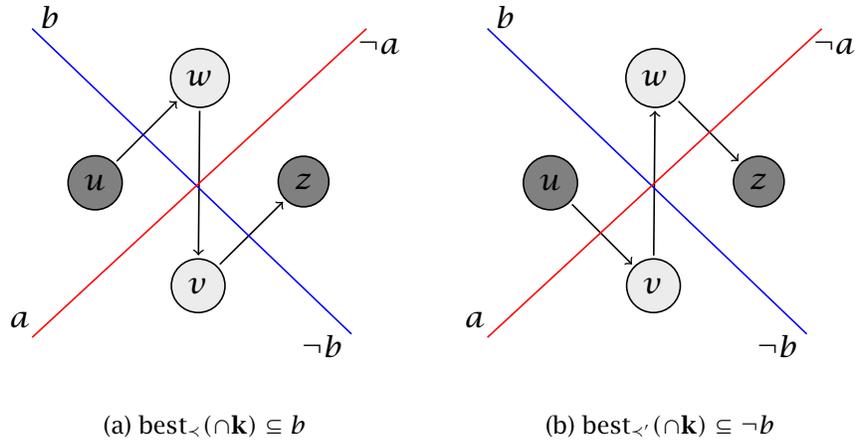


Figure 8 Faithfulness to $\mathbf{p} = \{a \parallel \top, \neg b \parallel \top\}$

faithful to this local preference, $\text{best}_{<}(a) \subseteq b$. And because $\mathbf{k} = \{a\}$ and $<$ was arbitrary it follows straight away that $\mathbf{s} \stackrel{|}{\sqsubseteq} b$.

Or take a predicament with conflict between local preferences like (6): you face the constraints that it's better to become mayor ($a \parallel \top$) and that it's better to avoid the pub ($\neg b \parallel \top$). Here, too, the two analyses agree on the bottom line: since you know that pub-going (b) is required for becoming mayor (a), it's not the case you ought to go to the pub and it's not the case that you ought not to stay away from the pub. But the way Analysis 2 reaches this conclusion is different.

To see this, consider any $<$ faithful to $\mathbf{p} = \{a \parallel \top, \neg b \parallel \top\}$. (For simplicity assume $\mathbf{A} = \{a, b\}$.) Every such ordering treats $a \cap \neg b$ as best simpliciter. Some of those orderings rank $a \cap b$ as next-best simpliciter and some rank $\neg a \cap \neg b$ as next-best simpliciter. Figure 6 shows two orderings faithful to \mathbf{p} exhibiting both features. Both orderings rank $a \cap \neg b$ as best simpliciter, but they differ on whether $a \cap b$ is next-best or whether $\neg a \cap \neg b$ is next-best: according to one $w < v$ and according to the other $v < w$. They can disagree about this relative goodness without either of them being unfaithful to any local preference in \mathbf{p} . Now, since in the predicament $\mathbf{k} = \{\neg a \cup b\}$, world u is the lone possibility ruled-out. Hence for some faithful $<$ we get that $\text{best}_{<}(\cap \mathbf{k})$ includes a b -world (w) and for some faithful $<$ we get that $\text{best}_{<}(\cap \mathbf{k})$ includes a $\neg b$ -world (v). And so $\mathbf{s} \not\stackrel{|}{\sqsubseteq} b$ and $\mathbf{s} \not\stackrel{|}{\sqsubseteq} \neg b$.

So far, so much convergence. A bit more: $\neg \Box$ isn't persistent.²⁸

Proposition 8. The sentence $\neg \Box b$ is not persistent in either \mathbf{k} or \mathbf{p} with respect to $\frac{\perp}{2}$.

This is where agreement ends. Take, for instance, predicaments like (8) in which one local preference (intuitively) over-rides another. Analysis 1 treats situations like this as cases of competition between local preferences. As a result, $\mathbf{s} \frac{\perp}{1} \neg \Box b$ (good) but also $\mathbf{s} \not\frac{\perp}{1} \Box \neg b$ (bad). The trouble is that in over-generating conflict, this analysis consults linearizations which are unfaithful to the local preferences.

This can't happen with Analysis 2: all and only faithful global preference orderings are relevant to *oughts*. So if $<$ is faithful to $\mathbf{p} = \{b \parallel \top, \neg b \parallel a\}$ then $\text{best}_{<}(a) \subseteq \neg b$. Since $\mathbf{k} = \{a\}$ it follows that $\mathbf{s} \frac{\perp}{2} \Box \neg b$. The over-riding local preference, well, over-rides the other one. This is a comparative good-making feature of this analysis.

To see that this isn't a vacuous prediction, we need to see that there is at least one $<$ faithful to \mathbf{p} . (There is.) That leads to the general idea of when a set of local preferences is consistent.

Definition 16 (Consistency). A set of local preferences \mathbf{p} is consistent iff there is a strict global ordering $<$ that is faithful to it.

This is the ordering analog of satisfiability. Using it we now know what it takes for a state to be consistent: there has to be a world compatible with \mathbf{k} and an ordering faithful to \mathbf{p} . In a state characterizing (8) the set of local preferences $\mathbf{p} = \{b \parallel \top, \neg b \parallel \top\}$ is consistent, as you'd expect and want. Of course, not just anything goes: $\{b \parallel a, \neg b \parallel a\}$ isn't consistent. This also goes in the pro-column.²⁹

²⁸ The limit cases show this. For non-persistence in \mathbf{k} consider $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \emptyset$ and $\mathbf{p} = \{b \parallel a\}$. Note that there are orderings faithful to \mathbf{p} that rank $\neg a \cap \neg b$ worlds as better than $a \cap b$ worlds and so $\mathbf{s} \frac{\perp}{2} \neg \Box b$. Contrast this with \mathbf{s}' that differs from \mathbf{s} only in that you know a . For non-persistence in \mathbf{p} consider $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a\}$ and $\mathbf{p} = \emptyset$ and then contrast with \mathbf{s}' that differs from this only in that you face the constraint $b \parallel a$.

²⁹ This is another spot where the difference between a preference-determination worldview (Analysis 1) and a preference-constraining worldview (Analysis 2) comes out: the mechanism for inducing a preference ordering always generates an ordering and so that mechanism treats what's going on in a set like $\{b \parallel a, \neg b \parallel a\}$ as a case of competition. Such sets seem more broken than that.

This can be turned into a characterization of what *oughts* are true in a predicament.³⁰

Theorem 2. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ be any state such that $\mathbf{s} \neq \perp$. Then $\mathbf{s} \models_2 \Box b$ iff $\mathbf{p} \cup \{\neg b \parallel \cap \mathbf{k}\}$ is inconsistent.

Analysis 2 is therefore in a precise sense exactly what we can get, prediction-wise, out of a set-up that deals with the indeterminacy of chunky preferences by quantifying over the global preference orderings that are faithful to them.

7 Merely possible constraints

The reason Analysis 2 smoothly handles predicaments like (8) in which one local preference over-rides another less specific one is that it uses the same engine that the flat (unembedded) fragment of the basic logic for variably strict conditionals uses.³¹ Famously, variably strict conditionals don't validate antecedent strengthening: given a nearness or similarity ordering, the nearest *a*-worlds can be *b*-worlds even though the nearest *a* \cap *c*-worlds aren't. In modeling predicaments using these tools we inherit this behavior. Therefore, from the fact that the best *a*-worlds are *b*-worlds in an ordering faithful to \mathbf{p} it doesn't follow that the best *a* \cap *c*-worlds must be, too. That is good news.

The bad news is that it is precisely this feature that dooms Analysis 2. Consider a predicament like (8) but in which you know just a little more, for instance that it is rainy:

(10) Better to go to the pub than not.	$b \parallel \top$
Better to not go to the pub given it's Sunday.	$\neg b \parallel a$
It is Sunday.	a
It is rainy.	c
You ought not to go to the pub.	$\mathbf{s} \models \Box \neg b$

Analysis 2 doesn't predict that you ought to stay away from the pub. The problem is simple: there are a lot of global orderings faithful to \mathbf{p} . In particular, there are some faithful to $\neg b \parallel a$ (a concrete preference you have) that are also faithful to $b \parallel a \cap c$ (a spectral preference you do not have).

³⁰ A probabilistic analog of this in the context of conditionals was first proved by Adams (1975).

³¹ See Burgess 1981, Veltman 1985. The flat fragment of the basic conditional logic also coincides (again, shared engine) with preferential entailment relations for non-monotonic logics; see Krauss et al. 1990, Makinson 1994.

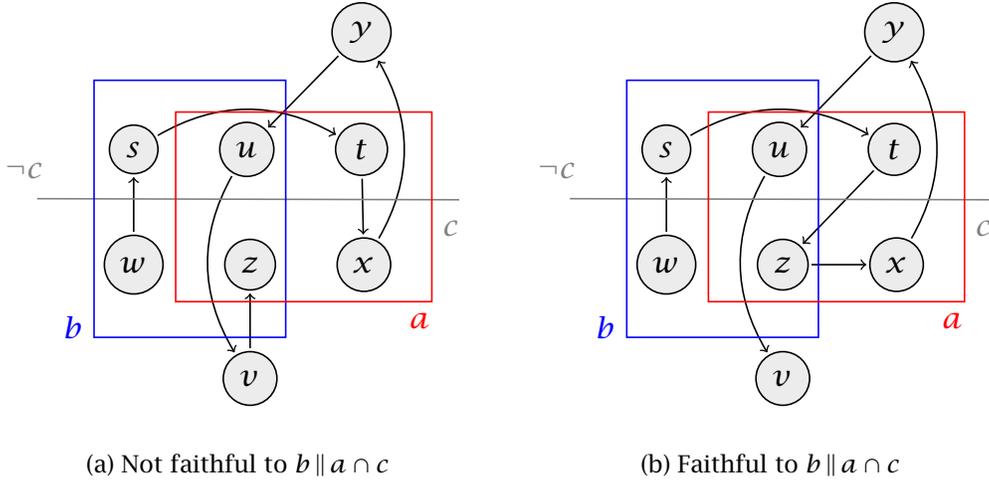


Figure 9 Some orderings faithful to $\mathbf{p} = \{b \parallel \top, \neg b \parallel a\}$

This, as we just saw, is the calling card of global orderings. Some of these orderings, in turn, are also faithful to the humble $b \parallel \top$. Let $<$ be one such witness. Now, any ordering like $<$ faithful to all three of these is also faithful to just $\mathbf{p} = \{b \parallel \top, \neg b \parallel a\}$, the local preferences in this predicament. So, when it comes to seeing whether $\mathbf{s} \stackrel{|}{\models} \square \neg b$ it follows that $<$ is among the orderings consulted. But since $\mathbf{k} = \{a, c\}$ that means we are concerned with the best $(a \cap c)$ -worlds according to $<$. This ordering is faithful to $b \parallel a \cap c$ and so $\text{best}_{<}(a \cap c) \subseteq b$. Therefore it can't be true that according to all orderings faithful to \mathbf{p} the best worlds compatible with \mathbf{k} are $\neg b$ -worlds and therefore $\mathbf{s} \not\stackrel{|}{\models} \square \neg b$. This isn't the right prediction.

To visualize things: Figure 9(a) shows an ordering that is faithful to both $b \parallel \top$ and $\neg b \parallel a$ but not faithful to the strengthened $b \parallel a \cap c$: the best $(a \cap c)$ -world according to it is x and that is a $\neg b$ -world. So far so good. However, the ordering in Figure 9(b) is also faithful to both $b \parallel \top$ and $\neg b \parallel a$. It is, in addition, also faithful to the merely possible $b \parallel a \cap c$: the best $(a \cap c)$ -world in this ranking is z and it is very much a b -world. If the constraints in your predicament are like \mathbf{p} then Analysis 2 quantifies over both sorts of orderings.

The prediction is bad, but the reason for it is worse. Your obligation to go to the pub has been over-ridden by the merely possible preference for avoiding the pub on rainy days together with the information that it is in fact

rainy. Rain, obviously, is not the issue. For any bit of unrelated information you have, there is the looming specter of a constraint you do not have that says: given that information, better to stay away from the pub. Any such merely possible local preference can undermine your obligation.

The general troublemaking feature is that the thing Analysis 2 leverages to get things right when one local preference seems to over-ride another is exactly the thing that makes $\Box b$ persistent in \mathbf{p} .

Corollary 3. For any b : $\Box b$ is persistent in \mathbf{p} with respect to $\frac{1}{2}$.

This certainly seems suboptimal: it threatens to make it too hard to have any obligations. Bt wait. If it is tough to find merely possible constraints that are consistent with a given \mathbf{p} maybe thngs aren't so bad. Alas, no, it is all too easy: exactly as easy as it is to indefinitely extend Sobel sequences.³² Thus, given Analysis 2, almost nothing is obligatory.

In fact, there is a constructive procedure for testing the consistency of a set of local preferences (we'll see it in the proof of Theorem 5). This can be used to find witnessing faithful global orderings that undermine almost any would-be obligation.³³

There is a natural metric on sets of local preferences.

Definition 17 (Openness). A set \mathbf{p} admits $b \parallel a$ iff there is a w such that:

- i. w confirms $b \parallel a$; and
- ii. w complies with every $b' \parallel a' \in \mathbf{p}$.

The set \mathbf{p} is open iff there is a $b \parallel a \in \mathbf{p}$ that it admits.

³² Like so:

- (i) a. If Alex comes, it will be fun;
- b. But if Alex and Billy come, it will be no fun;
- c. But if Alex and Billy and Chris come, it will be fun;
- d. ...

Lewis (1973) used the consistency of such sequences to argue that counterfactuals can't be any sort of strict conditional. (For dissenting views on that score see von Fintel 2001, Gillies 2007.)

³³ The construction was first used (again, in a quantitative setting) by Adams (1975). In the qualitative setting the procedure is part of system z (Pearl 1990); we'll see just how in Section 8.

In other words: there must be a $b \parallel a$ in \mathbf{p} and a w in $a \cap b$ such that w doesn't flout any member of \mathbf{p} .

Openness has a simple test, too.

Proposition 9. A set $\mathbf{p} = \{b_i \parallel a_i : 1 \leq i \leq n\}$ is open iff there is a $w \in a_1 \cup \dots \cup a_n$ such that w complies with every $b_i \parallel a_i \in \mathbf{p}$.

The promised result: \mathbf{p} is consistent iff every non-empty subset of \mathbf{p} is open. The proofs of the left-to-right and right-to-left directions (this one has the constructive procedure) are different enough that it makes sense to split them up.

Theorem 4. If $\mathbf{p} = \{b_i \parallel a_i : 1 \leq i \leq n\}$ is consistent then \mathbf{p} is open.

The converse claim (Theorem 5) is that every open set of local preferences is consistent. The proof makes use of two ideas. The first is a constructed ordered partition π of \mathbf{p} and the second is a ranking of worlds that reflects the priorities encoded in the partition. Before getting to the official definitions, a quick gloss on the pieces and how they'll fit together.

First, here's how to construct the ordered partition π of \mathbf{p} . Find the local preferences admitted by \mathbf{p} . Remove those; these are the first cell \mathbf{p}_0 of the partition. Next find the remaining local preferences that are admitted by $\mathbf{p} \setminus \mathbf{p}_0$. Remove those from $\mathbf{p} \setminus \mathbf{p}_0$; these form the second cell of the partition. Later, rinse, repeat until you run out of local preferences from the original \mathbf{p} .

An example: assume $\mathbf{A} = \{a, b\}$ and let $\mathbf{p} = \{b \parallel \top, \neg b \parallel a\}$. There is a $w \in b$ that complies with both local preferences (the $a \cap b$ -world). But $\neg b \parallel \top$ isn't admitted by \mathbf{p} : any world in $a \cap \neg b$ doesn't comply with $b \parallel \top$. So $\mathbf{p}_0 = \{b \parallel \top\}$. Looking at $\mathbf{p} \setminus \mathbf{p}_0$: there is a $v \in a \cap \neg b$ that complies with every local preference in this remainder so $\mathbf{p}_1 = \{\neg b \parallel a\}$. That is all of the local preferences so the partition π is $\langle \mathbf{p}_0, \mathbf{p}_1 \rangle$. Notice that the more specific preference occupies a later cell in the partition than does its less specific counterpart. That is key to how over-riding works.

And now the ranking. There is a natural ranking function that the ordered partition induces: it assigns higher ranks to (and thus penalizes) worlds that flout local preferences that occupy later cells in the partition. So, in our example, flouting $\neg b \parallel a$ is ceteris paribus worse than flouting $b \parallel \top$. To complete the proof sketch: every $<$ that (in a precise sense) tracks the ranking is faithful to \mathbf{p} and so \mathbf{p} is consistent.

Definition 18 (Ordered partitions). For any open set of local preferences \mathbf{p} , the ordered partition of \mathbf{p} is the partition $\pi = \langle \mathbf{p}_0, \dots, \mathbf{p}_n \rangle$ is defined inductively as follows:

- i. $\mathbf{p}_0 = \{b \parallel a : \mathbf{p} \text{ admits } b \parallel a\}$
- ii. $\mathbf{p}_{k+1} = \{b \parallel a \in \mathbf{p} \setminus \bigcup_{i=0}^k \mathbf{p}_i : \mathbf{p} \setminus \bigcup_{i=0}^k \mathbf{p}_i \text{ admits } b \parallel a\}$

Definition 19 (Induced ranking). Let $\pi = \langle \mathbf{p}_0, \dots, \mathbf{p}_n \rangle$ be the ordered partition of \mathbf{p} . The ranking function induced by π is the ranking $\kappa : W \rightarrow \mathbb{Z}_{\geq 0}$

$$\kappa_\pi(w) = \min \{i : w \text{ complies with } b \parallel a \text{ for every } b \parallel a \in \mathbf{p}_j, j \geq i\}$$

For any ranking κ , a strict global preference ordering represents κ iff w and v : if $\kappa(w) < \kappa(v)$ then $w < v$.

Theorem 5. If $\mathbf{p} = \{b_i \parallel a_i : 1 \leq i \leq n\}$ is open then \mathbf{p} is consistent.

The partition π is key. It sorts local preferences and thereby worlds that flout them. So, returning to the main issue: say you have the local preference $b \parallel a$. According to Analysis 2, $\Box b$ can be true only if it's impossible to pick up a local preference $\neg b \parallel a \cap c$ for any c that you happen to know. How widespread is the problem?

It's bad. Suppose $\mathbf{p} = \{b_i \parallel a_i : 1 \leq i \leq n\}$ is consistent but that there is a k such that $\mathbf{p}' = \mathbf{p} \cup \{\neg b_k \parallel a_k \cap c\}$ isn't. Since \mathbf{p} is consistent it is open (by Theorem 4) and so there is a w complies with every $b_i \parallel a_i \in \mathbf{p}$ (by Proposition 9). Since \mathbf{p}' is inconsistent every such w must flout $\neg b_k \parallel a_k \cap c$. That is: every witness to \mathbf{p} 's consistency must be in $a_k \cap c \cap b_k$. That's doable, sure, but hard.

For instance, take our example where $\mathbf{p} = \{b \parallel a\}$. Now suppose $a \cap b \not\subseteq c$ (not a high bar). Then $\mathbf{p}' = \mathbf{p} \cup \{\neg b \parallel a \cap c\}$ is consistent, too: every $(a \cap b \cap \neg c)$ -world is a witness to \mathbf{p} 's consistency but doesn't flout $\neg b \parallel a \cap c$. By Theorem 5 we can find an ordering by first constructing an ordered partition. Note that \mathbf{p}' admits $b \parallel a$ but does not admit $\neg b \parallel a \cap c$: any world confirming $\neg b \parallel a \cap c$ flouts $b \parallel a$ (Figure 7). So the partition is $\pi = \langle \mathbf{p}_0, \mathbf{p}_1 \rangle$ where $\mathbf{p}_0 = \{b \parallel a\}$ and $\mathbf{p}_1 = \{\neg b \parallel a \cap c\}$. This ensures that for any $<$ that represents this ranking, worlds flouting $\neg b \parallel a \cap c$ are worse than worlds that confirm it. Since any world confirming $\neg b \parallel a \cap c$ must flout $b \parallel a$, any of these orderings will do to undermine the obligation for Analysis 2.

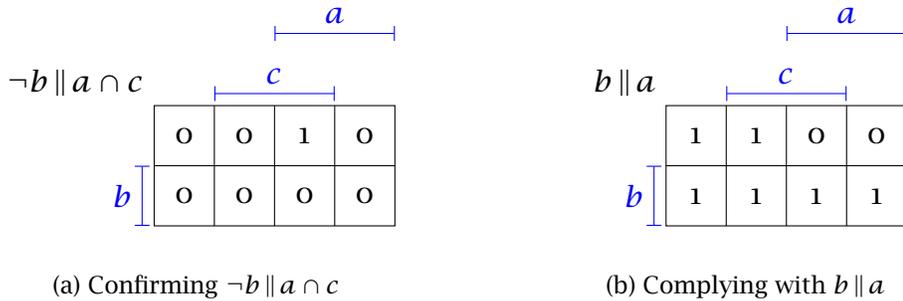


Figure 10 $\mathbf{p}' = \{b \parallel a, \neg b \parallel a \cap c\}$ doesn't admit $\neg b \parallel a \cap c$

Analysis 1 goes wrong by consulting orderings outside those faithful to a predicament. Analysis 2 instead says that all and only faithful orderings matter. This is why it can, and Analysis 1 can't, discriminate between conflicting constraints and over-riding constraints. But this is also exactly why the analysis goes wrong.³⁴

8 Constraining the constraints

Faithfulness allows too many global orderings through the front door and hence gets us too few obligations. Constraining things further is not easy, though. That is because we have been assuming throughout that the set \mathbf{S} of predicaments is non-trivial. This implies that it doesn't dictate what sorts of normative considerations you might face: since there are no a priori constraints on what local preferences you might come to have, there are likewise no a priori constraints on the space of global orderings.³⁵ An analysis that goes beyond Analysis 2 therefore has to put some boundaries in place and do this in a principled way.

The partition used in the proof of Theorem 5 is a way of enforcing boundaries and thus using more information by, in effect, ranking local preferences. Orderings sensitive to merely possible over-riding constraints,

³⁴ What might go for mushy credences thus can't go for chunky preferences. That is surprising since usually what goes for credence goes for preference and vice versa (you know, Ramsey (1929/1990) and so forth).

³⁵ This turns out to be enough to characterize the non-monotonic consequence relations built on global preference orderings: the basic preferential logic (a.k.a. the basic unembedded conditional logic) is complete with respect to such preferential models iff the space of orderings is rich in this way (see Halpern 2003).

while faithful, do not represent the ranking and so do not enter into things. Perhaps *ought* quantifies over only the best worlds compatible with what you know across this constrained set of orderings.³⁶

Analysis 3. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ be any state and π its ordered partition of \mathbf{p} .

$$\mathbf{s} \models_3 \Box b \text{ iff } \text{best}_<(\cap \mathbf{k}) \subseteq b$$

for every $<$ that represents κ_π .

This agrees with Analysis 2 as far as *ought* is concerned: whenever $\Box b$ is true in \mathbf{s} according to Analysis 2 then it is true in \mathbf{s} according to Analysis 3. No surprise there: \models_2 checks for truth according to all faithful orderings and each $<$ that represents κ_π is faithful to \mathbf{p} .

Proposition 10. For any state $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$: if $s \models_2 \Box b$ then $s \models_3 \Box b$.

This relationship is asymmetric when it comes to *oughts* (as opposed to their negations). The place to see this is in a concrete example like (10). Take it in stages. First, consider the predicament \mathbf{s} in which you have the constraints $\mathbf{p} = \{b \parallel \top, \neg b \parallel a\}$ and the information $\mathbf{k} = \{a \cap c\}$. The set \mathbf{p} admits $b \parallel \top$ but it doesn't admit $\neg b \parallel a$: no world confirming $\neg b \parallel a$ complies with $b \parallel \top$. So $\neg b \parallel a$ occupies a later cell in π and so in any $<$ that represents κ_π we know that $\text{best}_<(a \cap c) \subseteq \neg b$. Hence $\mathbf{s} \models_3 \Box \neg b$.

Now take the predicament \mathbf{s}' that extends this by having the constraint: $\mathbf{p}' = \mathbf{p} \cup \{b \parallel a \cap c\}$. The partition here extends the partition for \mathbf{p} : the last cell contains just $b \parallel a \cap c$ because it is not admitted by \mathbf{p} and it is not admitted by $\mathbf{p}_0 = \mathbf{p} \setminus \{b \parallel \top\}$ and it is not admitted by $\mathbf{p}_1 = \mathbf{p}_0 \setminus \{\neg b \parallel a\}$. So $b \parallel a \cap c$ takes precedence and in any $<$ that represents the induced ranking, $\text{best}_<(a \cap c) \subseteq b$. Hence $\mathbf{s}' \models_3 \Box b$. That gives us this:

Proposition 11. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a, c\}$ and $\mathbf{p} = \{b \parallel \top, \neg b \parallel a\}$. Consider the extension $\mathbf{s}' = \langle \mathbf{k}, \mathbf{p}' \rangle$ of \mathbf{s} where $\mathbf{p}' = \mathbf{p} \cup \{b \parallel a \cap c\}$. Then $\mathbf{s} \models_3 \Box \neg b$ but $\mathbf{s}' \not\models_3 \Box \neg b$. In fact, $\mathbf{s}' \models_3 \Box b$.

Corollary 6. The sentence $\Box b$ isn't persistent in either \mathbf{p} or \mathbf{k} with respect to \models_3 .

³⁶ This amounts, in the current framework, to system z entailment (Pearl 1990), which in turn is equivalent to what follows using rational consequence relations (see Lehmann & Magidor 1992, Makinson 1994).

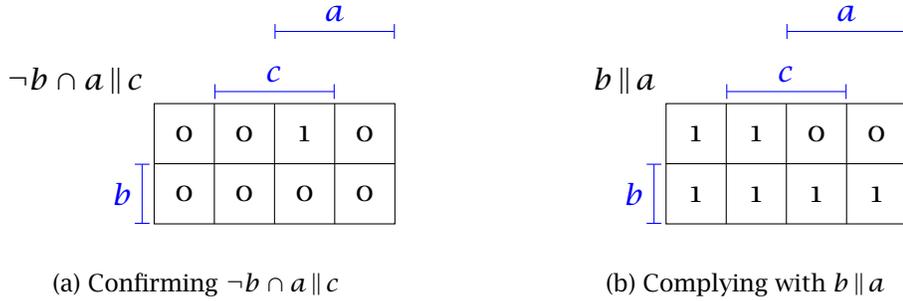


Figure 11 $\mathbf{p} = \{b \parallel a, \neg b \cap a \parallel c\}$ doesn't admit $\neg b \cap a \parallel c$

Analysis 3 does not discriminate between over-riding $b \parallel a$ by $\neg b \parallel a \cap c$ and over-riding it by $\neg b \cap a \parallel c$. Hence it treats (8) and (9) the same way. The set $\mathbf{p} = \{b \parallel a, \neg b \cap a \parallel c\}$ doesn't admit $\neg b \cap a \parallel c$. This is because any world confirming $\neg b \cap a \parallel c$ flouts $b \parallel a$. This is mapped in Figure 8, which looks a lot like the map in Figure 7. It should: $\neg b \parallel a \cap c$ and $\neg b \cap a \parallel c$ have the same confirming proposition.

While $\frac{\square}{3}$ extends $\frac{\square}{2}$ when it comes to $\square b$, it is both too strong and too weak. It is too strong in two ways. First: it has the wrong worldview. Analysis 3 tiptoes on the border of embracing the same local-preferences-determine-global-preferences stance that Analysis 1 takes. I can't shake the feeling that this denies the phenomenon of chunky preferences in the first place.

Second: Analysis 3 makes a priori judgments about which constraints carry more weight. It does this by ignoring some global preference orderings that are compatible with the local preferences being modeled. It has to in order to deliver more than Analysis 2. Among the ignored orderings are some that are faithful to merely possible constraints. It does this by minimizing the suboptimality represented in a given predicament.³⁷ Thus the ranking it is built on insists on a kind of normative equilibrium: no world in it can be made any better and stay faithful to the local preferences. We have no right to assume that.

³⁷ This basic idea is implemented in different ways in different nonmonotonic logics. See, for example, McCarthy 1980, Shoham 1987, Asher & Morreau 1991. If we are interested in modeling common sense or default reasoning, you can (maybe) whip up some enthusiasm for the idea: in that context it amounts to assuming that everything else (stuff not implicated by the defaults you have) is as normal as possible in every respect. Theorem 7 was first proved in Pearl 1990 (though in a slightly different way).

More precisely: in the induced κ_π , and so the orderings it determines, every world occupies the best position it possibly can.

Definition 20. Let κ be any ranking.

- i. κ is faithful to $b \parallel a$ iff every \prec that represents κ is faithful to $b \parallel a$.
- ii. κ is faithful to \mathbf{p} iff it is faithful to every $b \parallel a \in \mathbf{p}$.

Definition 21. For any ranking κ and any a : $\kappa(a) = \min \{\kappa(w) : w \in a\}$.

Lemma 1. For any ranking κ and local preference $b \parallel a$: κ is faithful to $b \parallel a$ iff $\kappa(a \cap b) < \kappa(a \cap \neg b)$.

Definition 22 (Improvements). Let κ be any ranking function and w any world. κ' is an improvement over κ iff for some w :

- i. $\kappa'(w) < \kappa(w)$; and
- ii. $\kappa'(v) = \kappa(v)$ for every $v \neq w$.

Such improvements are a limit case of Pareto-improving the situation in κ : κ' makes exactly one world better without disturbing the ranking of any other world. It's not always possible to do this in a way that is faithful to the local preferences of a predicament. One more lemma before seeing why.

Lemma 2. Let $\pi = \langle \mathbf{p}_0, \dots, \mathbf{p}_n \rangle$ be the ordered partition of \mathbf{p} and κ the associated ranking. If $b \parallel a \in \mathbf{p}_i$ then $\kappa(a \cap b) = i$.

Theorem 7. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ be any state and κ the ranking function associated with its ordered partition. If κ' is an improvement over κ then κ' is unfaithful to \mathbf{p} .

According to Analysis 3 every world is assumed to be as awesome as it can be modulo the local preferences. The world as we know it does not justify this optimism. Would that it were different.

Analysis 3 is also too weak. It imposes determinacy on predicaments by way of the ranking induced by the ordered partitions. Local preferences that are cellmates in that partition are thus treated on a par with each other. That's often enough fine but it can happen that in a predicament one of them doesn't really apply. In these cases Analysis 3 can't separate the chaff from

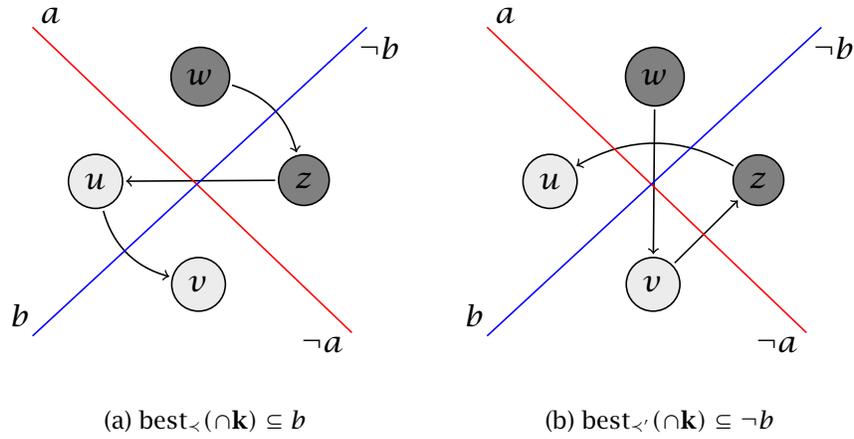


Figure 12 Representing κ_π

the independent wheat with the result that it under-generates obligations. There are two sorts of predicaments where this can happen.

The first sort is where you have multiple constraints but, as a matter of brute fact, complying with one of them is no longer possible. Suppose you face two constraints: that it is better to visit your aunt than not and that it is better to call your grandma than not. Ideally, of course, you ought to do both. But, pandemic, so you're not visiting your aunt.

- | | |
|--|--------------------------------|
| (11) It's better to go visit your aunt than not. | $a \parallel \top$ |
| It's better to call your grandma than not. | $b \parallel \top$ |
| You do not visit your aunt. | $\neg a$ |
| You ought to call your grandma. | $\mathbf{s} \models \square b$ |

An obligation to visit is not on the table but this should not interfere with an obligation to call.

But Analysis 3 disagrees. The set $\mathbf{p} = \{a \parallel \top, b \parallel \top\}$ admits both $a \parallel \top$ and $b \parallel \top$: any world in $a \cap b$ confirms both. They are thus cell-mates in the only cell of the ordered partition of \mathbf{p} . This means that worlds that flout one and confirm the other are ranked the same. That in turn means that if $<$ represents that ranking there is no guarantee that $\text{best}_{<}(\neg a) \subseteq b$. Two such orderings are pictured in Figure 8. According to one, $\text{best}_{<}(\cap \mathbf{k}) \subseteq b$ but according to another $\text{best}_{<}(\cap \mathbf{k}) \subseteq \neg b$.

Proposition 12. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{\neg a\}$ and $\mathbf{p} = \{a \parallel \top, b \parallel \top\}$. Then $\mathbf{s} \not\equiv_3 \Box b$.

Since $\mathbf{s} \equiv_2 \Box b$ only if $\mathbf{s} \equiv_3 \Box b$, these predicaments also show that Analysis 2 is too weak. The troublemaking $<$ that represents κ_π is, after all, faithful to \mathbf{p} . This sort of predicament is no trouble for Analysis 1: finding out that the optimal worlds where both a and b hold are out of reach because you know $\neg a$ does not get in the way of $\Box b$.

The second sort of predicament where Analysis 3 under-generates obligations involve situations in which you have multiple constraints with the same triggering information and (intuitively) just one of them gets over-ridden.

An example: you want to become mayor. This requires going to the pub and it requires chatting up the locals. But on Sundays, it would be a mistake to go to the pub.³⁸

(12) It's better to go to the pub.	$b \parallel \top$
It's better to chat up the locals.	$c \parallel \top$
It's better to not go to the pub, given it's Sunday.	$\neg b \parallel a$
It's Sunday.	a
You ought to not go to the pub.	$\mathbf{s} \models \Box \neg b$
You ought to chat up the locals.	$\mathbf{s} \models \Box c$

Neither of the unconditional local preferences $b \parallel \top$ and $c \parallel \top$ over-rides the other (\mathbf{p} admits them both) and so they are cellmates in \mathbf{p}_0 . Any world confirming $\neg b \parallel a$ flouts $b \parallel \top$, and so \mathbf{p} doesn't admit $\neg b \parallel a$. Thus $\neg b \parallel a$ occupies \mathbf{p}_1 . A world in $a \cap \neg b$ has to flout $b \parallel \top$, but nothing guarantees that such a world can't flout $c \parallel \top$, too. Some do, but since these local preferences are cellmates, those worlds are no worse than ones merely flouting $b \parallel \top$.

³⁸ That you have (unconditional) local preferences for pub-going and locals-chatting simplifies things by keeping the variables we have to track to a minimum. An example without this simplification:

(i) It's better to wear sunglasses, given a run.	$b \parallel a$
It's better to go early.	$c \parallel a$
Except in January: it's better to go later.	$\neg c \parallel a \cap d$
You're going out for a run in January.	$a \cap d$
You ought to wear sunglasses.	$\mathbf{s} \models \Box b$

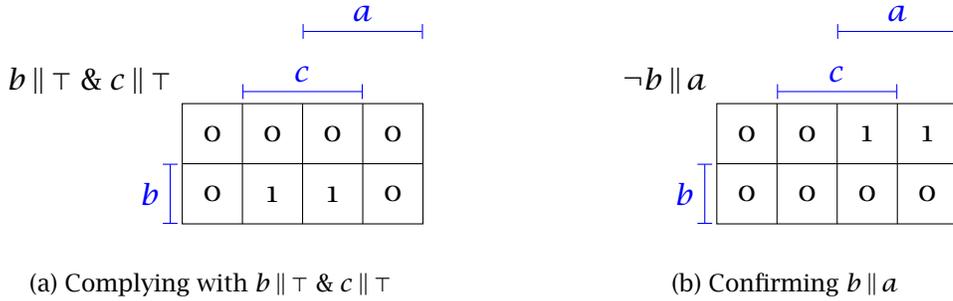


Figure 13 $\mathbf{p} = \{b \parallel \top, c \parallel \top, \neg b \parallel a\}$ doesn't admit $\neg b \parallel a$

There are thus two sorts of best a -worlds: those in $\neg b \cap c$ and those in $\neg b \cap \neg c$. The suboptimality of b has infected the erstwhile optimality of c . These dependencies that give rise to the ordered partition of \mathbf{p} are mapped in Figure 8.

Proposition 13. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a\}$ and $\mathbf{p} = \{b \parallel \top, c \parallel \top, \neg b \parallel a\}$. Then $\mathbf{s} \stackrel{|}{\equiv}_3 \square \neg b$ but $\mathbf{s} \not\stackrel{|}{\equiv}_3 \square c$.

As a corollary: these sorts of predicaments are equally problematic for Analysis 2.

They are also problematic for Analysis 1 but the reason is flipped around: according to Analysis 2 you ought to chat up the locals but it is not true that ought to stay away from the pub. There is both good here and bad. The mechanism for inducing $\preceq_{\mathbf{p}}$ treats $b \parallel \top$ and $c \parallel \top$ as independent, and that's good. But that same mechanism also treats $b \parallel \top$ and $\neg b \parallel a$ as independent, and that's bad. Within a the induced weak global ordering $\preceq_{\mathbf{p}}$ looks like Figure 14(a): worlds in $a \cap b \cap c$ are better than worlds in $a \cap b \cap \neg c$ and incomparable to worlds in $a \cap \neg b \cap c$, while worlds in $a \cap \neg b \cap c$ are better than worlds in $a \cap \neg b \cap \neg c$. One linearization of this has its best worlds within b (Figure 14(b)). It is unfaithful to $\neg b \parallel a$ and underwrites the mistaken prediction that \mathbf{s} doesn't support $\square \neg b$.

Proposition 14. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a\}$ and $\mathbf{p} = \{b \parallel \top, c \parallel \top, \neg b \parallel a\}$. Then $\mathbf{s} \stackrel{|}{\equiv}_1 \square c$ but $\mathbf{s} \not\stackrel{|}{\equiv}_1 \square \neg b$.

When it comes to obligation when you face independent constraints, some of which are trumped, there is a complementary distribution of wrong

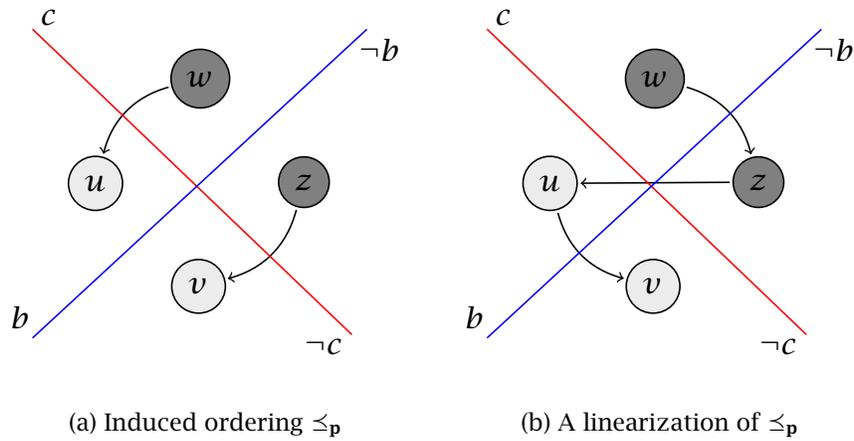


Figure 14 Too much independence within a

predictions between analyses that take local preferences to determine global preference orderings (like Analysis 1) and those that take local preferences to instead constrain admissible or relevant global preference orderings (like Analysis 2 and Analysis 3).

This same pattern holds if the over-riding is more subtle, similar to example (9).

- | | |
|--|----------------------------------|
| (13) Better to call before you go, given that you're going to visit. | $b \parallel a$ |
| Better to bring a gift, given that you're going to visit. | $c \parallel a$ |
| Better to not call but visit given there's a surprise party. | $\neg b \cap a \parallel d$ |
| You are going to visit and there's a surprise party. | $a \cap d$ |
| You ought to not call before. | $\mathbf{s} \models \Box \neg b$ |
| You ought to bring a gift. | $\mathbf{s} \models \Box c$ |

In the presence of the information that you are going to visit and there's a surprise party the local preference to call given that you're going shouldn't apply. Analysis 3 gets this right. The trouble is in doing so it also permits flouting of the local preference to bring a gift given you're going to visit.

To see this, map which local preferences are admitted by the set containing all of the local preferences in this predicament. Any world in $a \cap b \cap c$ confirms both $b \parallel a$ and $c \parallel a$. And since some of those are also in $\neg d$ (and so comply with $\neg b \cap a \parallel d$) both $b \parallel a$ and $c \parallel a$ are admitted. But $\neg b \cap a \parallel d$ isn't: Figure 15(a) maps the worlds confirming it while Figure 15(b) maps the

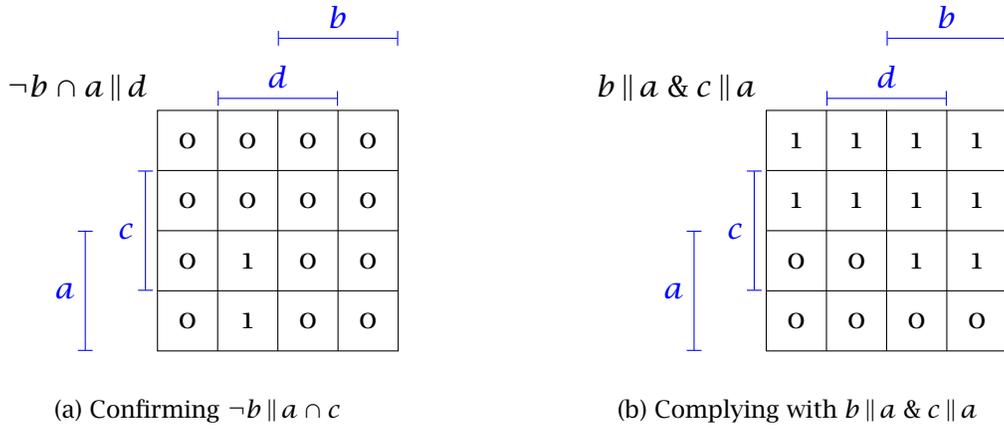


Figure 15 $\mathbf{p} = \{b \parallel a, c \parallel a, \neg b \cap a \parallel d\}$ doesn't admit $\neg b \cap a \parallel d$

worlds complying with both $b \parallel a$ and $c \parallel a$. Hence $\neg b \cap a \parallel d$ occupies a later cell in the ordered partition than $b \parallel a$ and $c \parallel a$ do. Thus when it comes to finding the best worlds compatible with $a \cap d$ in all the global orderings that represent this ranking, we get consensus about $\neg b$ ($\text{best}_{\prec}(a \cap d) \subseteq \neg b$ for every such \prec) but we don't get consensus about c (for some \prec it happens to be the case that $\text{best}_{\prec}(a \cap d) \subseteq c$ and for some $\text{best}_{\prec}(a \cap d) \not\subseteq c$). And so while $\mathbf{s} \models_3 \neg b$, unfortunately $\mathbf{s} \not\models_3 c$.

Proposition 15. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a\}$ and $\mathbf{p} = \{b \parallel a, c \parallel a, \neg b \cap a \parallel d\}$. Then $\mathbf{s} \models_3 \neg b$ but $\mathbf{s} \not\models_3 c$.

Again the troublemaking orderings that represent the induced ranking are faithful to \mathbf{p} and so equally troublemaking for Analysis 2. And again Analysis 1 reverses the order of under-generated obligations: according to it, you ought to bring a gift but it is not true that you ought to not call before.

Proposition 16. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a\}$ and $\mathbf{p} = \{b \parallel a, c \parallel a, \neg b \cap a \parallel d\}$. Then $\mathbf{s} \models_1 c$ but $\mathbf{s} \not\models_1 \neg b$.

This makes you wonder if, assuming the constraints imposed by the template that what you ought to do is what is best for a well-behaved relation of what-is-better-than-what, getting things right when one constraint overrides another is bought at the price of getting things wrong when you face multiple independent constraints and vice versa. Weighing in on that would

take this from a paper about a problem about preference to a paper about a solution to a problem about preference.

9 Tl;dr

I want to close with a more informal sketch of the argumentative arc we have been on.

We began with a template and an observation. The template says that *ought* is tightly tied to a relation of what-is-better-than-what: whether $\Box b$ is true in a predicament is a matter of *b*'s being true in the best worlds. What is best depends on some collection of one or more global orderings of better-than, where that collection is somehow tied to the the local preferences in the predicament. The observation is that normal predicaments under-determine such relations but don't seem to thereby under-determine the truth of *oughts*. Coping with this mismatch is the problem. Most of what we've been up to is showing that it is, in fact, a problem by showing that three candidate solutions (all natural implementations of familiar theories) can't work.

Analysis 1 is the simplest candidate. It says, first, that sets of local preferences determine global preference orderings in a straightforward way: *w* is at least as good (given a set of local preferences) as *v* iff *w* complies with every local preference that *v* does. It then says that $\Box b$ is true in a state iff the best worlds compatible with what you know are *b*-worlds in the induced weak global ordering. This candidate solution has good-making features: it predicts non-persistence in knowledge and in local preferences of both $\Box b$ and $\neg\Box b$. And it implements the standard-bearing architecture for relative modality.

These good-making features are bought at too great a price. The analysis denies the phenomenon of chunky preferences, saying that what seems like underdetermination really isn't. I think the worldview is mistaken and the way it is implemented by this analysis reveals reasons why. The mechanism for inducing global orderings from sets of local preferences incorrectly lumps different sorts of predicaments together: it can't tell the difference between cases of conflicting local preferences and cases of over-riding local preferences. In the relevant predicaments, Analysis 1 says we should consult global orderings that are not faithful to the local preferences we have. This leads to bad predictions, under-generating true *oughts*.

Analysis 2 adopts a different worldview: sets of local preferences don't determine global orderings, they constrain them. The constraining is simple: $\Box b$ is true iff the best worlds given what you know are all b -worlds in every global (strict, total) preference ordering faithful to your local preferences. So all and only faithful orderings count. This candidate solution, too, has good-making features: it agrees where it should with Analysis 1, and it rightly says that over-riding local preferences and conflicting local preferences are different sorts of things. It also leads to a clean characterization of what is obligated: $\Box b$ is true in a state $\langle \mathbf{k}, \mathbf{p} \rangle$ iff there is no ordering faithful to $\mathbf{p} \cup \{\neg b \parallel \cap \mathbf{k}\}$. Thus Analysis 2 is the *ought*-counterpart to the basic variably strict conditional logic. This is why this analysis gets things right for over-riding local preferences: variably strict conditionals don't go in for antecedent strengthening and that is just the shape that over-riding has.

However, this is also exactly why Analysis 2 wrongly predicts that *ought* is persistent in \mathbf{p} . Since \mathbf{p} underdetermines global orderings, there are a lot of orderings faithful to \mathbf{p} : if $b \parallel a$ is in \mathbf{p} then often enough we can find an ordering faithful to both $b \parallel a$ and in addition to the merely possible constraint $\neg b \parallel a \cap c$, because (again) no antecedent strengthening. If you happen to know $a \cap c$ then this merely possible constraint pre-empts your obligation to b . The result is a lot fewer true *oughts* than there should be.

So: Analysis 1 consults too few, and the wrong, global orderings and Analysis 2 consults way too many. Analysis 3 attempts to temper things: it doesn't consult any unfaithful global orderings but neither does it consult every faithful ordering. It achieves this by taking a wildly optimistic view about local preferences: that things are always as good as they can possibly be. This is too much. But it is also too little: the analysis still under-generates true *oughts*.

In particular, the way it makes sure that merely possible preferences don't get in the way also makes sure that it incorrectly lumps together worlds which flout one of your less important (over-ridden) constraints with worlds that flout all of your less important (not over-ridden) constraints. This is the mirror image of the problem facing Analysis 1: the way it makes sure that independent things stay independent makes sure that it wrongly over-generates conflict.

So the problem about preference is that we have a dilemma. On the one hand: predicaments don't seem to carry more information about the relative goodness of possibilities than is reflected in their sets of local preferences. But on the other hand: predicaments have to carry more information about

the relative goodness of possibilities than this — otherwise *oughts* wouldn't be determinately true in them. But just what this extra information is and where it comes from is a mystery.

Appendix: Proofs and proof-gestures

If you've made it this far, you deserve a cookie.³⁹ To accompany it, here are the proofs for the propositions, lemmas, and theorems from the main body of the paper.

Proposition 1. Both $\Box b$ and $\neg\Box b$ are non-persistent in \mathbf{k} with respect to \models_1 .

Proof. Non-persistence in \mathbf{k} of $\neg\Box b$: see examples (3) and (4). Non-persistence in \mathbf{k} of $\Box b$: see example (8). \square

Proposition 2. Both $\Box b$ and $\neg\Box b$ are non-persistent in \mathbf{p} with respect to \models_1 .

Proof. See examples (5) and (6). \square

Proposition 3. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a, c\}$ and $\mathbf{p} = \{b \parallel a, \neg b \parallel a \cap c\}$. Then $\mathbf{s} \not\models_1 \Box \neg b$.

Proof. See example (8). \square

Proposition 4. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a, c\}$ and $\mathbf{p} = \{b \parallel a, \neg b \cap a \parallel c\}$. Then $\mathbf{s} \not\models_1 \Box \neg b$.

Proof. There are two types of worlds in $\cap\mathbf{k}$: those in b and those in $\neg b$. Let w be any world in $a \cap b \cap c$ and v any world in $a \cap \neg b \cap c$. So w complies with $b \parallel a$ and flouts $\neg b \cap a \parallel c$ and v flouts $b \parallel a$ and complies with $\neg b \cap a \parallel c$. Thus $w \not\leq_{\mathbf{p}} v$ and $v \not\leq_{\mathbf{p}} w$. So $\text{best}_{\leq_{\mathbf{p}}}(\cap\mathbf{k})$ includes both w and v . \square

Theorem 1. Every weak global preference ordering has a linearization.

Proof. Consider any \leq on W where $|W| = n$. Let $\sigma: \wp(W) \rightarrow W$ be any choice function such that for any x , $\sigma(x) \in x$. Define a ranking function $\kappa: W \rightarrow \{0, \dots, n-1\}$ and a sequence of subsets x_0, \dots, x_{n-1} of W as follows. Let $x_0 = W$:

- i. For $\sigma(\text{best}_{\leq}(x_0)) = w$: $\kappa(w) = 0$ and $x_1 = x_0 \setminus \{w\}$;
- ii. For $\sigma(\text{best}_{\leq}(x_i)) = w$: $\kappa(w) = i$ and $x_{i+1} = x_i \setminus \{w\}$.

Note that κ is a bijection. Now let $w <' v$ iff $\kappa(w) < \kappa(v)$. The ordering $<'$ is a strict total global preference ordering: it is transitive and satisfies trichotomy simply because $<$ on the natural numbers does.

³⁹ <https://www.101cookbooks.com/triple-ginger-cookies>

So we need to see that \prec' agrees with the strict part of \preceq and resolves its incomparabilities. So suppose $w \prec v$. By construction there is a k such that $\{w, v\} \subseteq x_k$ but $\{w, v\} \not\subseteq x_{k+1}$. Hence either $\sigma(\text{best}_{\preceq}(x_k)) = w$ or $\sigma(\text{best}_{\preceq}(x_k)) = v$. Since $w \prec v$, it follows that $v \notin \text{best}_{\preceq}(x_k)$ and so $\sigma(\text{best}_{\preceq}(x_k)) \neq v$ and hence $\sigma(\text{best}_{\preceq}(x_k)) = w$. Thus: $\kappa(w) = k$ and $\kappa(v) > k$ and so $w \prec' v$. Finally: consider any w and v such that $w \not\prec v$ and $v \not\prec w$. If $w = v$ then $\kappa(w) = \kappa(v)$; if on the other hand $w \neq v$ then since κ is a bijection $\kappa(w) \neq \kappa(v)$ and so either $\kappa(w) < \kappa(v)$ or $\kappa(v) < \kappa(w)$ and so either $w \prec' v$ or $v \prec' w$. \square

Proposition 5. Fix a set \mathbf{p} of local preferences and its induced global ordering $\preceq_{\mathbf{p}}$. Let \prec_1, \dots, \prec_n be the linearizations of $\preceq_{\mathbf{p}}$. Then for any set x :

$$\text{best}_{\preceq_{\mathbf{p}}}(x) = \bigcup_{i=1}^n \text{best}_{\prec_i}(x)$$

Proof. Consider any $w \in \text{best}_{\preceq_{\mathbf{p}}}(x)$ and any $v \in x$ where $w \neq v$. It follows that $v \not\prec_{\mathbf{p}} w$. Hence there is a linearization \prec_i of $\preceq_{\mathbf{p}}$ such that $w \prec_i v$ and hence $v \not\prec_i w$. The choice of v was arbitrary so for no $v \in x$ is it the case that $w \neq v$ and $v \prec_i w$. Hence $w \in \text{best}_{\prec_i}(x)$ and so $w \in \bigcup_{i=1}^n \text{best}_{\prec_i}(x)$.

Going the other direction, consider any $w \notin \text{best}_{\preceq_{\mathbf{p}}}(x)$. We may further assume that $w \in x$ (otherwise it follows trivially that $w \notin \bigcup_{i=1}^n \text{best}_{\prec_i}(x)$). Since $w \in x$ but $w \notin \text{best}_{\preceq_{\mathbf{p}}}(x)$ there is a $v \in x$ such that $v \prec_{\mathbf{p}} w$. But since each \prec_i linearizes $\preceq_{\mathbf{p}}$ it follows that $v \prec_i w$ for each i . Hence for no i is it the case that $w \in \text{best}_{\prec_i}(x)$ and so $w \notin \bigcup_{i=1}^n \text{best}_{\prec_i}(x)$. \square

Proposition 6. A strict global preference ordering \prec is faithful to $b \parallel a$ iff $\text{best}_{\prec}(a) \subseteq b$.

Proof. Assume \prec is faithful to $b \parallel a$. Suppose for reductio that there is some $w \in \text{best}_{\prec}(a)$ such that $w \notin b$. Since \prec is faithful to $b \parallel a$ it follows that $a \cap b \prec a \setminus b$. Hence there is a $v \in a \cap b$ such that $v \prec w$ and so $w \notin \text{best}_{\prec}(a)$. Contradiction.

Now consider any $w \in a \setminus b$. We have to show that there is a $v \in (a \cap b)$ such that $v \prec w$. Consider $v \in \text{best}_{\prec}(a)$: since $w \in a$ it follows that $w \not\prec v$ and hence $v \prec w$ and so $a \cap b \preceq a \cap \neg b$. To see that $a \cap \neg b \not\preceq a \cap b$: suppose otherwise. Note that since $v \in \text{best}_{\prec}(a)$ and hence that $v \in a \cap b$ there would have to be a $u \in a$ such that $u \prec v$, contradicting the fact that $v \in \text{best}_{\prec}(a)$. Hence $a \cap b \prec a \cap \neg b$ and so \prec is faithful to $b \parallel a$. \square

Proposition 7. Let $\mathbf{p} = \{b \parallel \top, \neg b \parallel a\}$. There is a linearization of $\preceq_{\mathbf{p}}$ that is unfaithful to \mathbf{p} . Similarly, where $\mathbf{p}' = \{b \parallel a, \neg b \cap a \parallel c\}$, there is a linearization of $\preceq_{\mathbf{p}'}$ that is unfaithful to \mathbf{p}' . Moreover, where $\prec^x = \prec \cap (x \times x)$, there are linearizations of $\preceq_{\mathbf{p}'}$, \prec_1 and \prec_2 , such that \prec_1^a and \prec_2^c are unfaithful to \mathbf{p}' .

Proof. For $\mathbf{p} = \{b \parallel \top, \neg b \parallel a\}$, see example (8) and Figure 5. For $\mathbf{p}' = \{b \parallel a, \neg b \cap a \parallel c\}$ see example (9). In particular, since $z \prec_{\mathbf{p}'} w$ and $z \prec_{\mathbf{p}'} v$ (Figure 5), no linearization of $\preceq_{\mathbf{p}'}$ can be faithful to $\neg b \cap a \parallel c$. Figure 5(a) and Figure 5(b) represent two such linearizations of $\preceq_{\mathbf{p}'}$. Note that within c , the linearization \prec_1 is not faithful to $b \parallel a$ because $v \prec_1 w$. Similarly, within a , the linearization \prec_2 is not faithful to $\neg b \cap a \parallel c$ because $w \prec_2 v$. \square

Proposition 8. The sentence $\neg \Box b$ is not persistent in either \mathbf{k} or \mathbf{p} with respect to $\frac{\perp}{2}$.

Proof. Non-persistence in \mathbf{k} : let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \emptyset$ and $\mathbf{p} = \{b \parallel a\}$. Among the \prec faithful to \mathbf{p} : there is one in which $a \cap b$ is best full-stop and one in which $\neg a \cap \neg b$ is best full-stop. Hence $\mathbf{s} \models \neg \Box b$. Now consider $\mathbf{s}' = \langle \mathbf{k}', \mathbf{p}' \rangle$ where $\mathbf{k}' = \{a\}$ and $\mathbf{p}' = \mathbf{p}$. Let \prec be any ordering faithful to \mathbf{p}' . Hence: $\text{best}_{\prec}(a) \subseteq b$. Since $\cap \mathbf{k} = a$, it follows that $\mathbf{s}' \models \Box b$ and so $\mathbf{s}' \not\models \neg \Box b$.

Non-persistence in \mathbf{p} : let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a\}$ and $\mathbf{p} = \emptyset$. Clearly, $\mathbf{s} \models \neg \Box b$. Now consider $\mathbf{s}' = \langle \mathbf{k}', \mathbf{p}' \rangle$ where $\mathbf{k}' = \mathbf{k}$ and $\mathbf{p}' = \{b \parallel a\}$. Since \prec is faithful to \mathbf{p}' it follows that $\text{best}_{\prec}(a) \subseteq b$ and so since $\cap \mathbf{k}' = a$ that $\mathbf{s}' \models \Box b$ and hence $\mathbf{s}' \not\models \neg \Box b$. \square

Theorem 2. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ be any state such that $\mathbf{s} \neq \perp$. Then $\mathbf{s} \frac{\perp}{2} \Box b$ iff $\mathbf{p} \cup \{\neg b \parallel \cap \mathbf{k}\}$ is inconsistent.

Proof. Suppose (for reductio) that $\mathbf{s} \models \Box b$ but that $\mathbf{p} \cup \{\neg b \parallel \cap \mathbf{k}\}$ is consistent. Thus there is a \prec faithful to \mathbf{p} and faithful to $\neg b \parallel \cap \mathbf{k}$. Hence $\text{best}_{\prec}(\cap \mathbf{k}) \subseteq \neg b$. But $\mathbf{s} \models \Box b$ and so $\text{best}_{\prec}(\cap \mathbf{k}) \subseteq b$ and so $\cap \mathbf{k} = \emptyset$, contradicting the assumption that \mathbf{s} is consistent.

Now suppose that $\mathbf{p} \cup \{\neg b \parallel \cap \mathbf{k}\}$ is inconsistent. Consider any \prec faithful to \mathbf{p} . Since $\mathbf{p} \cup \{\neg b \parallel \cap \mathbf{k}\}$ is inconsistent clearly \prec can't be faithful to $\{\neg b \parallel \cap \mathbf{k}\}$. Hence $(\cap \mathbf{k} \cap \neg b) \not\subseteq (\cap \mathbf{k} \cap b)$ and so $(\cap \mathbf{k} \cap b) \prec (\cap \mathbf{k} \cap \neg b)$. Hence $\text{best}_{\prec}(\cap \mathbf{k}) \subseteq b$ and so $\mathbf{s} \models \Box b$. \square

Corollary 3. For any b : $\Box b$ is persistent in \mathbf{p} with respect to $\frac{\perp}{2}$.

Proof. This follows from Theorem 2: consider any (consistent) $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ and $\mathbf{s}' = \langle \mathbf{k}', \mathbf{p}' \rangle$ where $\mathbf{k} = \mathbf{k}'$ and $\mathbf{p} \subseteq \mathbf{p}'$. Suppose $\mathbf{s} \not\equiv_{\frac{1}{2}} \square b$. Hence $\mathbf{p} \cup \{\neg b \parallel \cap \mathbf{k}\}$ is inconsistent. But then so is any set that includes it and hence so is $\mathbf{p}' \cup \{\neg b \parallel \cap \mathbf{k}\}$. Thus $\mathbf{s}' \not\equiv_{\frac{1}{2}} \square b$. \square

Proposition 9. A set $\mathbf{p} = \{b_i \parallel a_i : 1 \leq i \leq n\}$ is open iff there is a $w \in a_1 \cup \dots \cup a_n$ such that w complies with every $b_i \parallel a_i \in \mathbf{p}$.

Proof. Suppose $\mathbf{p} = \{b_i \parallel a_i : 1 \leq i \leq n\}$ is open. So for some $b_i \parallel a_i \in \mathbf{p}$ and some $w: w \in (a_i \cap b_i)$ (and hence $w \in a_1 \cup \dots \cup a_n$) and w complies with every $b_j \parallel a_j \in \mathbf{p}$.

Suppose there is a $w \in a_1 \cup \dots \cup a_n$ such that w complies with every $b_i \parallel a_i \in \mathbf{p}$. Since w complies with every member of \mathbf{p} :

$$w \in ((\neg a_1 \cup b_1) \cap \dots \cap (\neg a_n \cup b_n))$$

And since $w \in a_1 \cup \dots \cup a_n$ it follows that $w \in a_i$ for some i . So: $w \in a_i$ and $w \in (\neg a_i \cup b_i)$ and so $w \in a_i \cap b_i$ for some i . Thus: \mathbf{p} admits $b_i \parallel a_i$ and is therefore open. \square

Theorem 5. If $\mathbf{p} = \{b_i \parallel a_i : 1 \leq i \leq n\}$ is open then \mathbf{p} is consistent.

Proof. Suppose (for reductio) that $\mathbf{p} = \{b_i \parallel a_i : 1 \leq i \leq n\}$ is consistent but not open. Let \prec be any ordering faithful to \mathbf{p} . Let $x = a_1 \cup \dots \cup a_n$. Since \mathbf{p} isn't open, for every $w \in x$ there is a $b_i \parallel a_i$ that w flouts. So take $w \in \text{best}_{\prec}(x)$: $w \in x$ and $w \prec v$ for every $v \in x$ such that $w \neq v$. So there is a $b_i \parallel a_i$ that w flouts: that is, $w \in a_i \cap \neg b_i$ (Proposition 9). Since \prec is faithful to \mathbf{p} and hence to $b_i \parallel a_i$ it follows that $a_i \cap b_i \prec a_i \cap \neg b_i$. Hence there is a $v \in a_i \cap b_i$ such that $v \prec w$. But $v \in x$ and so $v \not\prec w$, completing the reductio. \square

Theorem 5. If $\mathbf{p} = \{b_i \parallel a_i : 1 \leq i \leq n\}$ is open then \mathbf{p} is consistent.

Proof. Suppose every non-empty $\mathbf{p}' \subseteq \mathbf{p}$ is open. Let π be its ordered partition and κ_{π} the induced ranking. Note that for any w , $\kappa_{\pi}(w) = j$ iff $j = i + 1$ where i the largest index of partition cell that contains a local preference in \mathbf{p} that w flouts.

Let \prec_{π} be any ordering that represents κ_{π} . We show that for every i : if $b \parallel a \in \mathbf{p}_i$ then $\text{best}_{\prec_{\pi}}(a) \subseteq b$ and hence \prec_{π} is faithful to \mathbf{p} . So consider any $b \parallel a \in \mathbf{p}_i$ and suppose for reductio that there is a $w \in \text{best}_{\prec_{\pi}}(a)$ such that $w \notin b$. Hence $w \in a \cap \neg b$ and so w flouts $b \parallel a$ and so $\kappa_{\pi}(w) \geq i + 1$.

Now, let $\mathbf{p}^* = \mathbf{p} \setminus \bigcup_{k=0}^i \mathbf{p}_k$ and note that \mathbf{p}^* admits $b \parallel a$. So there is a $v \in a \cap b$ such that v complies with every member of \mathbf{p}^* . Hence $\kappa_\pi(v) \leq i$ and so $\kappa_\pi(v) < \kappa_\pi(w)$ and so $v \prec_\pi w$. But since $v \in a$ this contradicts the assumption that $w \in \text{best}_{\prec_\pi}(a)$, completing the proof. \square

Proposition 10. For any state $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$: if $s \stackrel{|}{=} \square b$ then $s \stackrel{|}{=} \square b$.

Proof. Consider any $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ such that $s \stackrel{|}{=} \square b$. So: $\text{best}_{\prec}(\cap \mathbf{k}) \subseteq b$ for every \prec faithful to \mathbf{p} . Consider any \prec' that represents the ordered partition of \mathbf{p} : \prec' is faithful to \mathbf{p} (Theorem 5) and so $\text{best}_{\prec'}(\cap \mathbf{k}) \subseteq b$ and hence $\mathbf{s} \stackrel{|}{=} \square b$. \square

Proposition 11. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a, c\}$ and $\mathbf{p} = \{b \parallel \top, \neg b \parallel a\}$. Consider the extension $\mathbf{s}' = \langle \mathbf{k}, \mathbf{p}' \rangle$ of \mathbf{s} where $\mathbf{p}' = \mathbf{p} \cup \{b \parallel a \cap c\}$. Then $\mathbf{s} \stackrel{|}{=} \square \neg b$ but $\mathbf{s}' \not\stackrel{|}{=} \square \neg b$. In fact, $\mathbf{s}' \stackrel{|}{=} \square b$.

Proof. The ordered partition π of \mathbf{p} is $\langle \mathbf{p}_0, \mathbf{p}_1 \rangle$ where $\mathbf{p}_0 = \{b \parallel \top\}$ and $\mathbf{p}_1 = \{\neg b \parallel a\}$. Consider any $v \in a \cap c$ such that $v \in b$. We will see that any $w \in a \cap c$ such that $w \in \neg b$ betters it. Since $v \in a \cap b$ it flouts $\neg b \parallel a$ and hence $\kappa_\pi(v) = 2$. Since $w \in a \cap \neg b$ it complies with $\neg b \parallel a$ and flouts $b \parallel \top$ and hence $\kappa_\pi(w) = 1$. So in every \prec that represents π : $w \prec v$. And so $\text{best}_{\prec}(a \cap c) \subseteq \neg b$.

Similarly, the partition π' of \mathbf{p}' is $\langle \mathbf{p}_0, \mathbf{p}_1, \mathbf{p}_2 \rangle$ where $\mathbf{p}_2 = \{b \parallel a \cap c\}$. Note that if $w \in (a \cap c) \cap \neg b$ then $\kappa_{\pi'}(w) = 3$: w flouts $b \parallel a \cap c$. And if $v \in (a \cap c) \cap b$ then $\kappa_{\pi'}(v) = 2$: v complies with $b \parallel a \cap c$ and flouts $\neg b \parallel a$. Hence in every \prec that represents π' : $v \prec w$ and so $\text{best}_{\prec}(a \cap c) \subseteq b$. \square

Corollary 6. The sentence $\square b$ isn't persistent in either \mathbf{p} or \mathbf{k} with respect to $\stackrel{|}{=} \square$.

Proof. Non-persistence in \mathbf{k} : compare $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ and $\mathbf{s}' = \langle \mathbf{k}', \mathbf{p} \rangle$ where $\mathbf{k} = \emptyset$, $\mathbf{k}' = \{a\}$, and $\mathbf{p} = \{b \parallel \top, \neg b \parallel a\}$. Non-persistence in \mathbf{p} : see Proposition 11. \square

Lemma 1. For any ranking κ and local preference $b \parallel a$: κ is faithful to $b \parallel a$ iff $\kappa(a \cap b) < \kappa(a \cap \neg b)$.

Proof. Suppose κ is faithful to $b \parallel a$ and consider any \prec that represents κ . Since κ is faithful to $b \parallel a$, $\kappa(a \cap b) < \kappa(a \cap \neg b)$. And since \prec represents κ , it follows that $a \cap b \prec a \cap \neg b$. So every $v \in a \cap \neg b$ is bettered by some

$w \in a \cap b$ and so $\min \{\kappa(w) : w \in a \cap b\} < \min \{\kappa(v) : w \in a \cap \neg b\}$ and so $\kappa(a \cap b) < \kappa(a \cap \neg b)$.

Now suppose $\kappa(a \cap b) < \kappa(a \cap \neg b)$ and consider any $v \in a \cap \neg b$. Let $<$ be any global preference ordering that represents κ . Note that $\kappa(a \cap \neg b) \leq \kappa(v)$. Now take a $w \in a \cap b$ such that $\kappa(w) = \kappa(a \cap b)$. Then we have:

$$\kappa(w) = \kappa(a \cap b) < \kappa(a \cap \neg b) \leq \kappa(v)$$

Thus: $\kappa(w) < \kappa(v)$. Since $<$ represents κ , $w < v$. So every $v \in a \cap \neg b$ is bettered by some $w \in a \cap b$ or other and so $a \cap b < a \cap \neg b$. \square

Lemma 2. Let $\pi = \langle \mathbf{p}_0, \dots, \mathbf{p}_n \rangle$ be the ordered partition of \mathbf{p} and κ the associated ranking. If $b \parallel a \in \mathbf{p}_i$ then $\kappa(a \cap b) = i$.

Proof. Suppose otherwise: so (i) $\kappa(a \cap b) > i$ or (ii) $\kappa(a \cap b) < i$. Suppose (i): since $b \parallel a \in \mathbf{p}_i$ there is a $v \in (a \cap b)$ such that $\kappa(v) = i$, in which case $\kappa(a \cap b) \not\leq i$. Suppose (ii): consider any $w \in \text{best}_\kappa(a \cap b)$ and let $\kappa(w) = j$ for some $j < i$. Since $\kappa(w) = j$, w complies with every member of \mathbf{p}_j . And since $w \in (a \cap b)$, \mathbf{p}_j thus admits $b \parallel a$ and hence it can't be in \mathbf{p}_i . Thus, since $\kappa(a \cap b) \not\leq i$ and $\kappa(a \cap b) \not\leq i$, $\kappa(a \cap b) = i$. \square

Theorem 7. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ be any state and κ the ranking function associated with its ordered partition. If κ' is an improvement over κ then κ' is unfaithful to \mathbf{p} .

Proof. Suppose κ' is an improvement over κ_π . So for some w and some i : $\kappa_\pi(w) = i + 1$ and $\kappa'(w) < i + 1$. Since $\kappa_\pi(w) = i + 1$ there is a $b \parallel a \in \mathbf{p}_i$ such that $w \in a \cap \neg b$. And since $b \parallel a \in \mathbf{p}_i$, it follows from Lemma 2 that $\kappa_\pi(a \cap b) = i$. Suppose for reductio that $\kappa'(a \cap b) < \kappa'(a \cap \neg b)$. Since $w \notin a \cap b$, $\kappa_\pi(a \cap b) = \kappa'(a \cap b) = i$. And so $\kappa'(a \cap \neg b) > i$. But since $w \in a \cap \neg b$ and $\kappa'(w) < i + 1$, we also have $\kappa'(a \cap \neg b) < i + 1$. Hence: $\kappa'(a \cap b) \not\leq \kappa'(a \cap \neg b)$ and so by Lemma 1 κ' isn't faithful to $b \parallel a$ and hence unfaithful to \mathbf{p} . \square

Proposition 12. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{\neg a\}$ and $\mathbf{p} = \{a \parallel \top, b \parallel \top\}$. Then $\mathbf{s} \not\equiv_3 \square b$.

Proof. Let $W = \{u, v, w, z\}$ be as in Figure 8. The ordered partition of \mathbf{p} is $\pi = \langle \mathbf{p}_0 \rangle$ where $\mathbf{p}_0 = \mathbf{p}$. Hence $\kappa_\pi(x) = 0$ if $x = w$ and $\kappa_\pi(x) = 1$ if $x \neq w$. There are three global preference orderings that represent κ_π ; there are some $<$ such that $\text{best}_<(\cap \mathbf{k}) \subseteq b$ (Figure 12(a)) and some $<'$ such that $\text{best}_{<' }(\cap \mathbf{k}) \subseteq \neg b$ (Figure 12(b)). \square

Proposition 13. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a\}$ and $\mathbf{p} = \{b \parallel \top, c \parallel \top, \neg b \parallel a\}$. Then $\mathbf{s} \stackrel{\parallel}{\equiv}_3 \square \neg b$ but $\mathbf{s} \not\stackrel{\parallel}{\equiv}_3 \square c$.

Proof. The set \mathbf{p} admits both $b \parallel \top$ and $b \parallel \top$ but doesn't admit $\neg b \parallel a$ and so the ordered partition is $\pi = \langle \mathbf{p}_0, \mathbf{p}_1 \rangle$ where $\mathbf{p}_0 = \{b \parallel \top, c \parallel \top\}$ and $\mathbf{p}_1 = \{\neg b \parallel a\}$. (See Figure 8.) Given this partition:

$$\begin{aligned} \kappa_\pi(w) &= 0 & \text{for every } w &\in \neg a \cap b \cap c \\ \kappa_\pi(w) &= 1 & \text{for every } w &\in a \cap \neg b \\ \kappa_\pi(w) &= 2 & \text{for every } w &\in a \cap b \end{aligned}$$

Consider any \prec that represents κ_π and any $w \in \text{best}_\prec(a)$. Notice that $\kappa_\pi(w) \neq 0$ (because then $w \in \neg a$) and $\kappa_\pi(w) \neq 2$ (because there are $v \in a \cap \neg b$ such that $\kappa_\pi(v) < \kappa_\pi(w)$ and so $v \prec w$ and hence w wouldn't be in $\text{best}_\prec(a)$). Hence $\kappa_\pi(w) = 1$. Hence $w \in a \cap \neg b$ and so $\text{best}(\cap \mathbf{k}) \subseteq \neg b$ and thus $\mathbf{s} \stackrel{\parallel}{\equiv}_3 \square \neg b$.

Some members of $a \cap \neg b$ are in c and some aren't. So consider any $v \in a \cap \neg b \cap c$ and $v' \in a \cap \neg b \cap \neg c$. Since $\kappa_\pi(v) = \kappa_\pi(v')$ there is some \prec representing κ_π such that $v' \prec v$. Notice that $\text{best}_\prec(a) = a \cap b \cap \neg c$. So $\text{best}_\prec(\cap \mathbf{k}) \not\subseteq c$ and so $\mathbf{s} \not\stackrel{\parallel}{\equiv}_3 \square c$. \square

Proposition 14. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a\}$ and $\mathbf{p} = \{b \parallel \top, c \parallel \top, \neg b \parallel a\}$. Then $\mathbf{s} \stackrel{\parallel}{\equiv}_1 \square c$ but $\mathbf{s} \not\stackrel{\parallel}{\equiv}_1 \square \neg b$.

Proof. Any world in $a \cap b \cap c$ only complies with $b \parallel \top$ and $c \parallel \top$ while any world in $a \cap \neg b \cap c$ only complies with $c \parallel \top$ and $\neg b \parallel a$. Hence these worlds are incomparable. Worlds in $a \cap b \cap \neg c$ comply only with $b \parallel \top$ and worlds in $a \cap \neg b \cap \neg c$ comply only with $c \parallel \top$. Hence both of these sorts of worlds are bettered by both of the first sorts. Hence $\text{best}_\mathbf{p}(a) \subset c$ but not $\text{best}_\mathbf{p}(a) \subseteq \neg b$. \square

Proposition 15. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a\}$ and $\mathbf{p} = \{b \parallel a, c \parallel a, \neg b \cap a \parallel d\}$. Then $\mathbf{s} \stackrel{\parallel}{\equiv}_3 \square \neg b$ but $\mathbf{s} \not\stackrel{\parallel}{\equiv}_3 \square c$.

Proof. The proof is largely like the proof of Proposition 13. \square

Proposition 16. Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a\}$ and $\mathbf{p} = \{b \parallel a, c \parallel a, \neg b \cap a \parallel d\}$. Then $\mathbf{s} \stackrel{\parallel}{\equiv}_1 \square c$ but $\mathbf{s} \not\stackrel{\parallel}{\equiv}_1 \square \neg b$.

Proof. The proof (sketch) is largely like the proof of Proposition 14. \square

References

- Adams, Ernest W. 1975. *The Logic of Conditionals*. D. Reidel.
- Asher, Nicholas & Michael Morreau. 1991. Commonsense entailment: A modal theory of nonmonotonic reasoning. In *Jelia '90*, 387–392. European Workshop on Logics in AI, Springer-Verlag.
- Burgess, John P. 1981. Quick completeness proofs for some logics of conditionals. *Notre Dame Journal of Formal Logic* 22(1). 76–84.
- Cariani, Fabrizio, Stephan Kaufmann & Magdalena Kaufmann. 2013. Deliberative modality under epistemic uncertainty. *Linguistics and Philosophy* 36. 225–259. <http://dx.doi.org/10.1007/s10988-013-9134-4>.
- Etherington, David W. & Raymond Reiter. 1983. On inheritance hierarchies with exceptions. In *Proceedings of the third national conference on artificial intelligence*, 104–108.
- Fine, Kit. 1975. Vagueness, truth and logic. *Synthese* 30(3–4). 265–300.
- von Fintel, Kai. 2001. Counterfactuals in a dynamic context. In Michael Kenstowicz (ed.), *Ken Hale: A life in language*. MIT Press.
- von Fintel, Kai. 2012. The best we can (expect to) get? challenges to the classic semantics for deontic modals. URL <http://mit.edu/fintel/fintel-2012-apa-ought.pdf>. Ms, MIT.
- 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. Springer.
- van Fraassen, Bas C. 1974. Hidden variables in conditional logic. *Theoria* 40(3). 176–190.
- Gillies, Anthony S. 2007. Counterfactual scorekeeping. *Linguistics and Philosophy* 30. 329–360. <http://dx.doi.org/10.1007/s10988-007-90186>.
- Gillies, Anthony S. 2014. Default logic and *Hafta*. URL <https://semanticsarchive.net/Archive/TdjOGIoZ/>. Author Meets Critics on Horty's *Reasons as Defaults*.
- Halpern, Joseph Y. 2003. *Reasoning about uncertainty*. MIT Press.
- Horty, John. 2014. Deontic modals: Why abandon the classical semantics? *Pacific Philosophical Quarterly* 95. 424–460. <http://dx.doi.org/10.1111/papq.12044>.
- Horty, John F. 2012. *Reasons as defaults*. Oxford University Press.
- Joyce, Jim. 2011. A defense of imprecise credences in inference and decision making. In Tamar Gendler Szabó & John Hawthorne (eds.), *Oxford studies in epistemology*, vol. 4. Oxford University Press.

- Kolodny, Niko & John MacFarlane. 2010. Ifs and oughts. *The Journal of Philosophy* 107(3). 115-143.
- Kratzer, Angelika. 1981a. The notional category of modality. In H. J. Eikmeyer & H. Rieser (eds.), *Words, worlds, and contexts. new approaches in word semantics*, 38-74. Berlin: de Gruyter.
- Kratzer, Angelika. 1981b. Partition and revision: The semantics of counterfactuals. *Journal of Philosophical Logic* 10. 201-16.
- Kratzer, Angelika. 1991. Modality. In Arnim von Stechow & Dieter Wunderlich (eds.), *Semantics: An international handbook of contemporary research*, 639-650. Berlin: de Gruyter.
- Kratzer, Angelika. 2012. *Modals and conditionals*. Oxford University Press.
- Krauss, Sarit, Daniel Lehmann & Menachem Magidor. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44(1-2). 167-207.
- Lehmann, Daniel & Menachem Magidor. 1992. What does a conditional knowledge base entail? *Artificial Intelligence* 55(1). 1-60.
- Lewis, David. 1973. *Counterfactuals*. Blackwell.
- Lewis, David. 1981. Ordering semantics and premise semantics for counterfactuals. *Journal of Philosophical Logic* 10. 217-234.
- Makinson, David. 1994. General patterns in nonmonotonic reasoning. In Dov Gabbay, C. J. Hogger & J. A. Robinson (eds.), *Handbook of logic in artificial intelligence and logic programming*, vol. 3, 35-110. Oxford University Press.
- McCarthy, John. 1980. Circumscription: A form of non-monotonic reasoning. *Artificial Intelligence* 13. 27-39.
- Pearl, Judea. 1990. System Z: A natural ordering of defaults with tractable applications to default reasoning proceedings third conference on theoretical aspects of reasoning about knowledge, pacific grove, ca (1990) 121-135. In Moshe Y. Vardi (ed.), *Proceedings of tark*, 121-135. Morgan Kaufman.
- Pollock, John L. 1976. The 'possible worlds' analysis of counterfactuals. *Philosophical Studies* 29(6). 469-476.
- Pollock, John L. 1998. The logical foundations of goal-regression planning in autonomous agents. *Artificial Intelligence* 106. 267-334.
- Ramsey, Frank P. 1929/1990. General propositions and causality. In H. A. Mellor (ed.), *F. Ramsey, Philosophical Papers*. Cambridge University Press.
- Reiter, Raymond & Giovanni Criscuolo. 1981. On interacting defaults. In *Proceedings of the seventh international joint conference on artificial intel-*

- ligence*, 270-276. Morgan Kaufman.
- Shoham, Yoav. 1987. A semantical approach to nonmonotonic logics. In *Proceedings of the 2nd symposium on logic in computer science*, 275-279.
- Spohn, Wolfgang. 1988. Ordinal Conditional Functions: A Dynamic Theory of Epistemic States. In William Harper & Brian Skyrms (eds.), *Causation in Decision, Belief Change, and Statistics*, vol. 2, 105-134. Reidel.
- Spohn, Wolfgang. forthcoming. Defeasible moral reasoning. *Synthese* <http://dx.doi.org/10.1007/s11229-019-02083-2>.
- Stalnaker, Robert. 1968. A theory of conditionals. In Nicholas Rescher (ed.), *Studies in logical theory*, 98-112. Basil Blackwell.
- Stalnaker, Robert & Richmond Thomason. 1970. A semantic analysis of conditional logic. *Theoria* 36. 23-42.
- Szpilrajn, Edward. 1930. Sur l'extension de l'ordre partiel. *Fundamenta Mathematica* 16(1). 386-389.
- Veltman, Frank. 1976. Prejudices, presuppositions and the theory of counterfactuals. *Amsterdam Papers in Formal Grammar: Proceedings of the 1st Amsterdam Colloquium* 248-281.
- Veltman, Frank. 1985. *Logics for conditionals*. Ph.D. thesis, University of Amsterdam.