# A problem about preference*

Anthony S. Gillies
*Rutgers University*

## 1 Nutshell (introductory)

Obligation describing language (and so, I will assume, obligation) is hooked up with preference, a relation of what-is-better-than-what. But ordinary situations 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 trouble-making and that getting out of said 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 become 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, 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, 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 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 *ought*s of whatever flavor so long as it is tied to a relation of comparative betterness

(1) a. I ought do the dishes.
(given the house rules plus what we know)

b. Jimbo ought to be in class.
(given the terms of his scholarship plus what we know)

c. Lisa ought to go to the rally.
(given her promises plus what we know)

d. Homer ought to move his car.
(given the laws plus what we know)

We can't do what we know can't be.[2] We are obligated to do what's true in the best of the possibly suboptimal possibilities compatible with what we know.

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 situation says that the best worlds, with respect to the relevant relation of what-is-better-than-what in that situation and given what we know in it, are *b* worlds. That gives us a general template (we'll use □ for the target modal):

**Template.** □*b* is true in a situation iff all the best worlds given the normative constraints plus what we know in that situation are *b*-worlds.

Our job is to fill this in by saying how the normative constraints 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. The problem is that 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

---

of that same flavor. I try to stay away from examples with full-on moral *ought*s so that substantive debates about moral betterness don't distract from the structural point.

2 Not all occurrences of deontic modals are like that. For instance:

(2) 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 prejacents. These are perfectly good *ought*s, but not our topic.

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. $\mathbf{L}^{\mathbf{o}}_{\mathbf{A}}$ is the smallest set containing $\mathbf{A}$ that is closed under $\neg$ and $\cap$.

ii. $\mathbf{L}_{\mathbf{A}}$ is the smallest set including $\mathbf{L}^{\mathbf{o}}_{\mathbf{A}}$ and is such that: if $a \in \mathbf{L}^{\mathbf{o}}_{\mathbf{A}}$ then $\Box a \in \mathbf{L}_{\mathbf{A}}$; and if $a \in \mathbf{L}_{\mathbf{A}}$ then $\neg a \in \mathbf{L}_{\mathbf{A}}$.

iii. $w$ is a possible world iff $w \in 2^{\mathbf{A}}$.

For readability I will be willfully sloppy and let $a, b, c, \ldots$ range both over non-modal sentences and the sets of worlds where the sentences in question are true. 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).[3]

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 better than $\neg b$. Write it this way: $b \parallel a$. So $b \parallel a$ is a claim in the metalanguage saying that there is 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 situation. In that case: you have a local preference, within the you-promised worlds, for doing the review rather than not doing it.

**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 basic formal properties local preferences must have. I will put these in terms of minimal constraints on the space of consistent states (and thereby the predicaments they model).

3 Similarly for $\cup$. When it doesn't make much difference, let's also suppress mention of $\mathbf{A}$ and $\mathbf{L}_{\mathbf{A}}$ even though, officially, everything is parametric on a choice of underlying language.

**Postulate 1.** $\mathbf{s}$ is consistent 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.

**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 descriptive $a$:

i. $\mathbf{s} \models a$ iff $\cap \mathbf{k} \subseteq a$;

ii. $\mathbf{s} \models \neg \Box a$ iff $\mathbf{s} \not\models \Box a$.

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 are obliged to do $a$ or not. So while this definition doesn't yet say what it takes for sentences like $\Box b$ to be supported or true in a state, it does put constraints on possible analyses by requiring that they don't permit obligation gaps. 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.[4]

4 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.

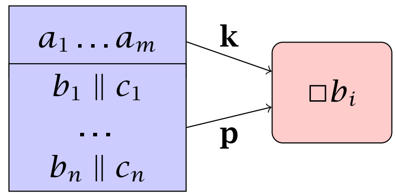


**Figure 1** Predicament: **k** (set of propositions) + **p** (set of local preferences)

Now to the problem. Take the ordinary predicament above: your only relevant knowledge is that you promised to do the review ($a$) and the the only normative constraint is that it is 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.[5] Still, there are *ought*s that are clearly true.

| | | |
|---|---|---|
| (3) | Better to review (than not), given you promised. | $b \parallel a$ |
| | You promised. | $a$ |
| | You ought do the review. | $\mathbf{s} \models \Box b$ |

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.[6]

5 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 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.

6 These need not (yet) be classical, economist-approved preference relations since we are (for now) leaving open the possibility that $\preceq$ is not connected: there may be a $w$ and $v$ such that $w \npreceq v$ and $v \npreceq w$.

**Definition 6** (Global preference ordering)**.** A (weak) preference ordering $\preceq$ is a reflexive and transitive relation over the set $W$ of worlds; $\prec$ is the strict part of $\preceq$: $w \prec v$ iff $w \preceq v$ but $v \not\preceq w$.

The best worlds in a set $x$ with resect to $\preceq$ are those in $x$ that are not dominated by any others in $x$.

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

$$\text{best}_{\preceq}(x) = \{w \in x\colon \text{ there is no } v \in x \text{ s.t. } v \prec w\}$$

Such global orderings (seem to) contain a lot more information about comparative goodness than your run of the mill set of local preferences do. Still, since given such an ordering it is straightforward to find the best worlds in a set with respect to it, the hope is global preference orders can be leveraged to represent the information in a set of local preferences, bridging the gap between local preferences and the *ought*s that are true based on them.

## 4 Preference determination

Sets of local preferences are 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.** 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 $b \parallel a$.

The set $(a \cap \neg b)$ is the flouting proposition for $b \parallel a$ and $(\neg a \cup b)$ is its complying proposition.

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

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

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.[7] What we have done is to treat local preferences as the ingredients of premise sets or, as they are now more generally known, ordering sources for deontic modals.

This naturally pairs with the template.

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

$$\mathbf{s} \models_1 \Box b \text{ iff } \text{best}_{\preceq_{\mathbf{p}}}(\cap \mathbf{k}) \subseteq b$$

It is a two step procedure: derive orderings from the set of complying propositions for $\mathbf{p}$ (an ordering source); *ought* quantifies over the resulting best worlds compatible with the set of propositions characterizing what is known (a modal base).[8]

This solution has a lot going for it. It rightly predicts that *ought*s 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.[9] Then you are unconstrained: you can do whatever, pubwise.

| | | |
|---|---|---|
| (4) | 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.

| | | |
|---|---|---|
| (5) | 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. | $\mathbf{s} \models \Box b$ |

7 Lewis 1981 argued that the difference between premise semantics and ordering semantics for counterfactuals is exaggerated: every set of premises (propositions) induces an ordering in this way that is in a precise sense equivalent as far as evaluating counterfactuals is concerned (see, e.g., Gillies 2017: §4,8).

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

9 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.

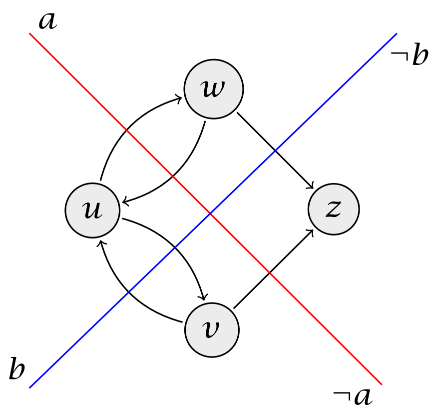


(a) No promise

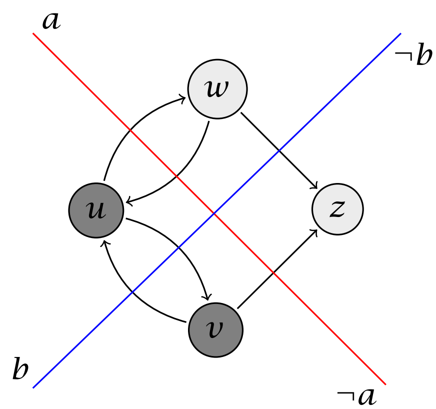


(b) Promise

**Figure 2** Promise keeping

Analysis 1 predicts this pattern smoothly. Here's why (this is all pictured in Figure 2(a)).[10] In the case of (4) the lone local preference induces an ordering $\preceq_{\mathbf{p}}$ that divides the worlds into two clumps: those that verify the preference (i.e., those in $\neg a \cup b$) and those that flout it (i.e., those in $a \cap \neg b$). Each clump is an equivalence class with respect to $\preceq_{\mathbf{p}}$ 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 $\text{best}_{\preceq_{\mathbf{p}}}(\cap\mathbf{k})$ 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 (5). 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.

10 Conventions for the graphs: $w \preceq v$ 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.

**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.

**Proposition 1.** $\Box b$ 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 (4) and (5). (I will put off discussing non-persistence in $\mathbf{k}$ of $\Box b$ until Section 5). □

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.

| | | |
|---|---|---|
| (6) | 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.[11]

| | | |
|---|---|---|
| (7) | 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. | $\mathbf{s}' \not\models \Box\neg b$ |

11 The example (from Kratzer 1981) is an instance of what is called the "Nixon diamond" in nonmonotonic/default logic circles. 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?

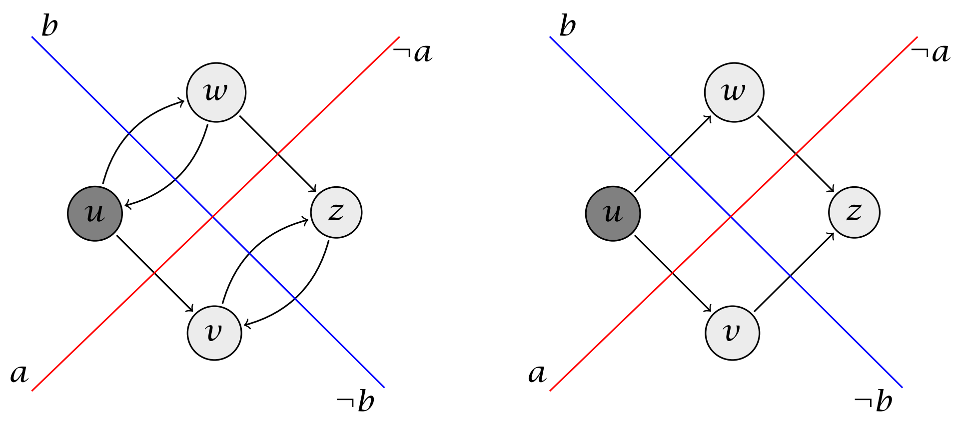


(a) No competition

(b) Competition

**Figure 3** Becoming mayor

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.[12]

Again, Analysis 1 predicts this pattern smoothly. When the only local preference is $a \parallel \top$ then as far as $\preceq_{\mathbf{p}}$ 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 good (or bad, as it happens) as each other. In the ordering in Figure 3(a), $u$ is the only world ruled out by **k**. There is a path from $w$ to $v$ but not vice versa, so $\text{best}_{\mathbf{p}}(\cap \mathbf{k}) = \{w\}$ and so $\mathbf{s} \models_1 \Box b$.

When you instead have the local preferences $\mathbf{p}' = \{a \parallel \top, \neg b \parallel \top\}$, worlds that were previously tied with respect to $\preceq_{\mathbf{p}}$ 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 $v$ any more. 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 **k**

12 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.

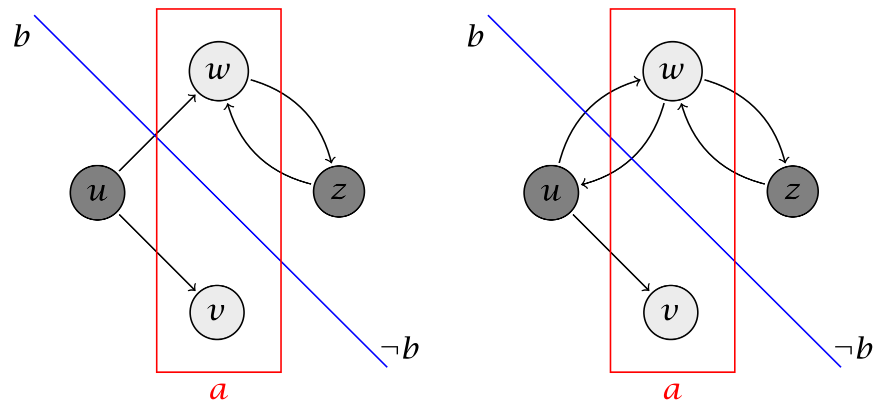


(a) Using both local preferences (b) Using one local preference

**Figure 4** Specificity

according to $\preceq_{\mathbf{p}}$ are $w$ and $v$. Since one is a $b$-world and one a $\neg b$-world it follows that $\mathbf{s}' \not\models_1 \Box b$ and $\mathbf{s}' \not\models_1 \Box\neg b$.

The general pattern:

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

*Proof.* See examples (6) and (7). □

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

## 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.

The mechanism that induces the orderings enforces categorical priorities. Each world's relative position in the induced global preference ordering is determined by comparing which categorical priorities each complies with. This way of adjudicating trade-offs between local preferences overgenerates incomparabilities between worlds.

Predictably, this has empirical consequences. An example:

(8) 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. $a$
You ought not go to the pub. $\mathbf{s} \models \Box\neg b$

The judgment that you ought not go to the pub is not shared by Analysis 1. It agrees that it is not the case that you ought to go (this is the promised non-persistence in $\mathbf{k}$ of $\Box b$ from Proposition 1), but it stops short of saying that you ought to not go.

The official and slightly more general form of the problem:

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

*Proof.* Figure 4(a) contains a simple model of (8). Here $u$ is the only world complying with both local preferences, $w$ and $v$ each comply with one but not the other, and $z$ flouts them both. So $w$ and $v$ are incomparable in the induced $\preceq_{\mathbf{p}}$. Since $a = \{w, v\}$, both are in $\text{best}_{\preceq_{\mathbf{p}}}(\cap\mathbf{k})$. But $w \in \neg b$ and $v \in b$ and so both $\mathbf{s} \not\models_1 \Box b$ and $\mathbf{s} \not\models_1 \Box\neg b$. □

This doesn't square with the judgment that $\Box\neg b$ is true in (8). The local preference to stay away from the pub on Sundays should over-ride the local preference for pub going; instead we get an incomparability between worlds complying with one and not the other of the local preferences.[13]

A possible patch: 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 $\neg b \parallel a \cap c$. Since $\top$ is weaker than $a$, in modeling (8) we should look at $\mathbf{p}' = \{\neg b \parallel a\}$. A model using the coarser $\preceq_{\mathbf{p}'}$ is in Figure 4(b). It is true that $\text{best}_{\preceq_{\mathbf{p}'}}(\cap\mathbf{k}) = \{w\}$ and so this patch may seem to do the job.

But, no. First, it is plainly ad hoc. The patch requires that we, qua theorists, are doing all the weighing and adjudicating between local preferences. This is work that the mechanism of inducing an ordering was supposed to do. Second, there is no simple a priori story for what counts as "more specific": it can depend on what you know not just on what local preferences you have.

To see this, consider a somewhat richer state where $\mathbf{k} = \{a, a \subseteq c\}$ and $\mathbf{p} = \{b \parallel a, \neg b \parallel c\}$. For concreteness suppose you face the constraint that it is better to go to the pub given that Alex goes and you know it is a brute fact

13 The incomparability in $\preceq_{\mathbf{p}}$ is not the problem: insisting instead on indifference between $w$ and $v$ (by appeal to some as yet specified departure from Definition 9) doesn't help.

that Chris will be at the pub only if Alex is and that Chris will be there. Add to this the constraint that it is better to not go given that Chris will be there. Given what you know, this is a more specific, over-riding consideration.

(9) Better to go to the pub given Alex is going. $b \parallel a$
Better to not go to the pub given Chris is going. $\neg b \parallel c$
Chris goes. $c$
Either Chris doesn't go or Alex goes. $\neg c \cup a$
You ought not go to the pub. $\mathbf{s} \models \Box\neg b$

This is not a case of conflict: you should not go to the pub. But neither local preference has a triggering condition that is asymmetrically stronger than the other's. Figure 4(a) can be interpreted as a partial representation of a model of this state using the induced preference ordering $\preceq_{\mathbf{p}}$: pictured are just the $c$-worlds. Here, as in the original model of (8), $w$ and $v$ are incomparable but are both in $\text{best}_{\preceq_{\mathbf{p}}}(\cap\mathbf{k})$ and so *ought* quantifies over both and that clashes with the judgment that $\Box\neg b$ is supported. Since what is more specific isn't merely a matter of what local preferences you have, the patch doesn't work.

The mechanism for inducing global preference orderings lumps specificity situations like (8) together with situations like (7) that have genuinely competing local preferences. The lumping is wrong, but the reason why it so lumps is more wrong.[14] To get at that, we'll be concerned with strict preference relations.

**Definition 11.** A strict global preference relation $\prec$ is a transitive relation on $W$ such that for any $w, v$ exactly of one of the following holds: $w \prec v$, $w = v$, or $v \prec w$.

These are economist-approved strict preference relations.[15]

**Definition 12** (Linearizations)**.** Let $\preceq$ be any global preference order (reflexive and transitive relation over $W$). A strict preference $\prec'$ linearizes $\preceq$ iff:

14 The problem is not unique to Analysis 1. The problem lies with the way 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. 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.

15 The second condition in the definition is sometimes called "trichotomy". Given such a $\prec$, define its weak counterpart as follows: for any $w$ and $v$, $w \preceq v$ iff $w \prec v$ or $w = v$. The resulting relation is a weak total ordering of $W$.

i. if $w \prec v$ then $w \prec' v$; and

ii. if $w \nprec v$ and $v \nprec w$ then either $w \prec' v$ or $v \prec' w$.

If $\prec'$ linearizes $\preceq$ it is a linearization of $\preceq$.

In general an induced $\preceq_{\mathbf{p}}$ has a bunch of linearizations: one for each way of settling each tie or incomparability in $\preceq_{\mathbf{p}}$.[16]

**Proposition 4.** Fix a set **p** of local preferences and its induced global ordering $\preceq_{\mathbf{p}}$. Let $\prec_1, \ldots, \prec_n$ 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 $v \npreceq_{\mathbf{p}} w$. Hence there is a linearization $\prec_i$ of $\preceq_{\mathbf{p}}$ such that $w \prec_i v$ and hence $v \nprec_i w$. 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 $w \notin \bigcup_{i=1}^{n} \text{best}_{\prec_i}(x)$. □

We therefore lose nothing by talking about a set of linearizations of $\preceq_{\mathbf{p}}$ rather than the induced $\preceq_{\mathbf{p}}$ itself. As far as Analysis 1 is concerned it doesn't matter which.

So, in (8), we know that $\mathbf{s} \not\models_1 \Box\neg b$. Thinking in terms of the linearizations of $\preceq_{\mathbf{p}}$: this is because there is a linearizing $\prec$ according to which the best worlds in $a$ are in $b$ but also a linearizing $\prec'$ according to which the best worlds in $a$ are in $\neg b$ (Figure 5). It is not merely that some linearizations go beyond the normative information in **p**. Some of them are in a precise sense unfaithful to it. To say just what that amounts to, first lift the strict global ranking $\prec$ on worlds to one on propositions.

16 This idea, too, has roots in conditionals: in particular, in Lewis's (1981) argument showing that there isn't much disagreement between versions of ordering semantics that allow ties and incomparabilities (Pollock 1976) and those that allow ties but no incomparabilities (Lewis 1973) and those that allow neither (Stalnaker 1968, Stalnaker & Thomason 1970). Linearizations also play a role in partial-order planning (see, for instance, Pollock 1998). Note that applying $\text{best}_{\prec}(\cdot)$ where $\prec$ is a strict global preference will always return a singleton.

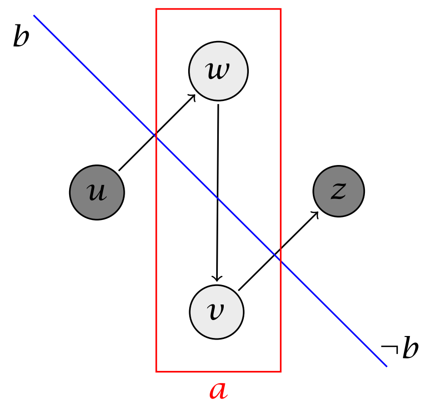


(a) $\text{best}_{\prec}(\cap \mathbf{k}) \subseteq \neg b$

(b) $\text{best}_{\prec'}(\cap \mathbf{k}) \subseteq b$

**Figure 5** Two linearizations of $\preceq_{\mathbf{p}}$

**Definition 13** (Lifted preferences)**.** For any strict preference ordering $\prec$ on $W$ it's 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 $b \not\preceq^+ a$.

So $a$ is better than $b$ just when every $b \cap \neg a$ world is dominated 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.[17]

**Definition 14** (Faithfulness)**.** A strict global preference ordering $\prec$ is faithful to $b \parallel a$ iff $a \cap b \prec a \setminus b$; 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 as better than the flouting proposition. This is different from the induced ordering mechanism.

**Proposition 5.** A strict global preference ordering $\prec$ is faithful to $b \parallel a$ iff $\text{best}_{\prec}(a) \subseteq b$.

17 Since it is generally clear from context whether $\prec$ or $\prec^+$ is in play, from now on I'll omit the superscript.

*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 $w \nprec v$ and hence $v \prec w$ and so $a \cap b \preceq a \cap \neg b$. To see that $a \cap \neg b \npreceq a \cap b$: 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 $b \parallel a$. □

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. And so our obligation describing language can end up appealing to global orderings that, intuitively, it shouldn't. That is the case in predicaments like (8) where one local preference intuitively over-rides a more specific one. Officially:

**Proposition 6.** Let $\mathbf{p} = \{b \parallel \top, \neg b \parallel a\}$. There is a linearization of $\preceq_{\mathbf{p}}$ that is unfaithful to $\mathbf{p}$.

*Proof.* See example (8) and Figure 5. □

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 ends 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: 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 such relations. What counts as relevant? Faithfulness.[18]

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

$$\mathbf{s} \models_2 \Box b \text{ iff } \text{best}_{\prec}(\cap \mathbf{k}) \subseteq b$$

for every $\prec$ 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: $\models_2$ treats predicaments as constraining (and not in general determining) the global orderings relevant for *ought*s.

This gets a lot right. It embraces the idea that local preferences underdetermine a proper global preference relation. But it copes with it gracefully, agreeing with a lot of the verdicts that Analysis 1 gets right.

Take, for starters, the predicament in (5): 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. Here's why. Consider any $\prec$ faithful to $\mathbf{p}$. Then $\text{best}_{\prec}(a) \subseteq b$. Since $\mathbf{k} = \{a\}$ and $\prec$ was arbitrary it follows straight away that $\mathbf{s} \models_2 \Box b$.

Or take a predicament with conflict between local preferences like (7). Here, too, the two analyses agree on the bottom line: it's not the case you ought to go to the pub and it's not the case that you ought not to go to the pub. But the way we get to this conclusion by Analysis 2 is different.

To see this, consider any $\prec$ faithful to $\mathbf{p} = \{a \parallel \top, \neg b \parallel \top\}$. Every such ordering treats $a \cap \neg b$ worlds as best simpliciter. Some of those orderings rank $a \cap b$ worlds as next-best simpliciter and some rank $\neg a \cap \neg b$ as next-best simpliciter. Figure **??** shows two orderings faithful to $\mathbf{p}$ exhibiting both

18 In addition to its conditionals roots, there is also a clear connection from supervaluationism (Fine 1975). It also looks like the preference-analog of modeling mushy credences with sets of probability functions: there the suggestion (e.g. in Joyce 2011) is that the epistemic states of rational agents should be modeled by sets of coherent probability functions and that such agents should update their credences by conditionalizing on those sets. These family resemblances make a lot of sense since the idea is that predicaments are shot through with a specific kind of indeterminacy.

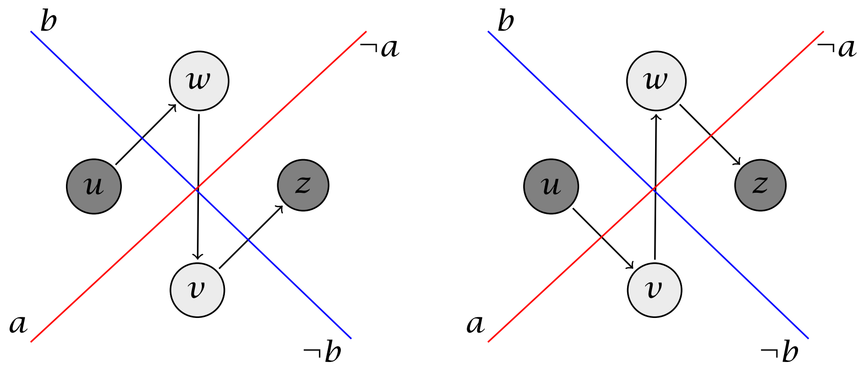


(a) $\text{best}_{\prec}(\cap\mathbf{k}) \subseteq b$    (b) $\text{best}_{\prec'}(\cap\mathbf{k}) \subseteq \neg b$

**Figure 6** Faithfulness to $\preceq_{\mathbf{p}} = \{a \,\|\, \top, \neg b \,\|\, \top\}$

features where $a = \{w, u\}$ and $b = \{w, z\}$. Both orderings rank $a \cap \neg b$ worlds (i.e., $u$) 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 $w \prec v$ and according to the other $v \prec w$. This divergence of opinion doesn't make either of them unfaithful to $\mathbf{p}$. Now, since in the predicament $\mathbf{k} = \{\neg a \cup b\}$, world $u$ is the lone possibility ruled-out. Hence for some faithful $\prec$ we get that $\text{best}_{\prec}(\cap\mathbf{k})$ includes a $b$-world ($w$) and for some faithful $\prec$ we get that $\text{best}_{\prec}(\cap\mathbf{k})$ includes a $\neg b$-world ($v$). And so $\mathbf{s} \not\models_2 \Box b$ and $\mathbf{s} \not\models_2 \Box\neg b$.

So far, so much convergence. A bit more:

**Proposition 7.** $\neg\Box b$ is not persistent in either $\mathbf{k}$ or $\mathbf{p}$ with respect to $\models_2$.

*Proof.* Non-persistence in $\mathbf{k}$: let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \emptyset$ and $\mathbf{p} = \{b \,\|\, a\}$. 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 ay 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 $\mathbf{p}' = \{b \,\|\, a\}$. 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 $\mathbf{s}' \not\models \neg\Box b$. □

This is where agreement ends. Take, for instance, predicaments like (8). Analysis 1 treats situations like this as cases of competition between local preferences. As a result, $\mathbf{s} \models_1 \neg\Box b$ (good) but also $\mathbf{s} \not\models_1 \Box\neg b$ (bad). Not so Analysis 2. The reason is straightforward: if $\prec$ is faithful to $\mathbf{p} = \{b \parallel \top, \neg b \parallel a\}$ then $\text{best}_\prec(a) \subseteq \neg b$. Since $\mathbf{k} = \{a\}$ it follows that $\mathbf{s} \models_2 \Box\neg b$. The over-riding local preference, well, over-rides the other one. This is a relevant comparative good-making feature.

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

**Definition 15** (Consistency)**.** A set of local preferences $\mathbf{p}$ is consistent iff there is a strict global ordering $\prec$ 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. So $\mathbf{p} = \{a \parallel \top, \neg b \parallel \top\}$ is consistent, as we'd expect and want. Of course, not just anything goes: $\{b \parallel a, \neg b \parallel a\}$ isn't, by these lights, consistent. This also goes in the pro-column.[19]

This can be turned into a characterization of what *ought*s are true in a predicament.[20]

**Theorem 1.** For any consistent state $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$: $\mathbf{s} \models_2 \Box b$ 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 $(\cap\mathbf{k} \cap \neg b) \not\prec (\cap\mathbf{k} \cap b)$ 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 $\mathbf{s} \models \Box b$. □

19 This is another spot where the difference in between a preference-determination worldview (Analysis 1) and a preference-constraining worldview (Analysis 2) comes out: inducing preferences in the usual way 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.

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

Analysis 2 is therefore in a precise sense exactly what we can get, prediction-wise, out of a set-up that takes deals with the indeterminacy of chunky preferences by quantifying over the global preference rankings 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 exactly the same engine that the analog of the flat (unembedded) fragment of the basic logic for variably strict conditionals uses.[21] 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, in any global ordering faithful to a set of normative constraints, just because the best $a$-worlds are $b$-worlds 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, (say) 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 go to the pub. $\mathbf{s} \models \Box\neg b$

The judgment that you should stay away from the pub is not shared by $\models_2$. Here's why: there are a lot of global orderings. In particular, there are some faithful to both $\neg b \parallel a$ and $b \parallel a \cap c$. This is the calling card of global preference orderings. Some of these, in turn, are also faithful to the humble $b \parallel \top$. Let $\prec$ be one such witness. Now, any ordering like $\prec$ 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 $\mathbf{s} \models_2 \Box\neg b$ it follows that $\prec$ is among the orderings consulted. But since $\mathbf{k} = \{a, c\}$ that means

21 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.

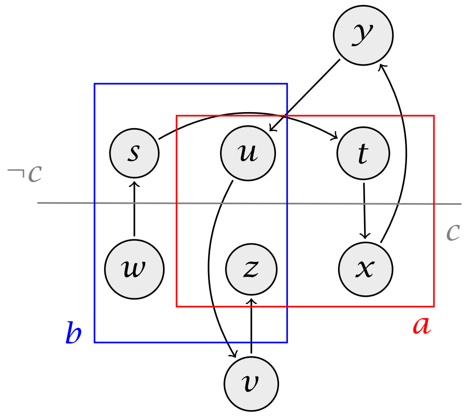


(a) Not faithful to $b \parallel a \cap c$

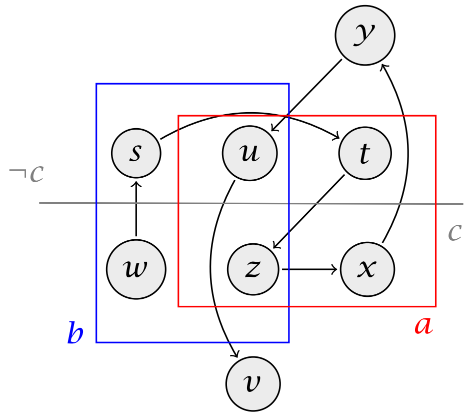


(b) Faithful to $b \parallel a \cap c$

**Figure 7** Some orderings faithful to $\mathbf{p} = \{b \parallel \top, \neg b \parallel a\}$

we are concerned with the best $(a \cap c)$-worlds according to $\prec$. This ordering is faithful to $b \parallel a \cap c$ and so $\text{best}_{\prec}(\cap \mathbf{k}) \subseteq b$. 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 $\mathbf{s} \not\models_2 \neg b$. This isn't the right prediction.

A concrete example: Figure 7(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 global ordering in Figure 7(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 your state is 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 actual information that it is rainy. Rain is not the issue: for any bit of unrelated information you have, there is the looming specter that it is a priori possible to face a constraint that says: given that information, better to stay away from the pub.

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 it too hard to have obligations: $\Box b$ is guaranteed to be persistent in $\mathbf{p}$.

**Corollary 2.** $\Box b$ is persistent in $\mathbf{p}$ with respect to $\models_2$.

*Proof.* This follows from Theorem 1: consider any (consistent) $\mathbf{s} = \langle \mathbf{k}, \mathbf{s} \rangle$ and $\mathbf{s}' = \langle \mathbf{k}', \mathbf{s}' \rangle$ where $\mathbf{k} = \mathbf{k}'$ and $\mathbf{p} \subseteq \mathbf{p}'$. Suppose $\mathbf{s} \models_2 \Box b$. Hence $\mathbf{p} \cup \{\neg b \,\|\, \cap \mathbf{k}\}$ is inconsistent. But then so is any set that includes it and hence so is $\mathbf{p}' \cup \{\neg b \,\|\, \cap \mathbf{k}\}$. Thus $\mathbf{s}' \models_2 \Box b$. □

This certainly seems suboptimal. Is there room for hope, though? So far we don't know exactly how easy it is to find merely possible constraints that are consistent with a given $\mathbf{p}$. Alas, it is all too easy: exactly as easy as it is to indefinitely extend Sobel sequences. Thus, given Analysis 2, almost nothing is obligatory.[22]

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 4). This can be leveraged to finding witnessing faithful global orderings that undermine almost any would-be obligation.[23]

**Definition 16** (Admissibility)**.** A set $\mathbf{p}$ admits $b \,\|\, a$ iff there is a $w$ such that:

i. $w \in a \cap b$; and

ii. $w$ complies with every $b' \,\|\, a' \in \mathbf{p}$.

The set $\mathbf{p}$ is admissible iff there is a $b \,\|\, a \in \mathbf{p}$ that it admits.

The admissibility of $\mathbf{p}$ has a simple test, too.

**Proposition 8.** A set $\mathbf{p} = \{b_i \,\|\, a_i \colon 1 \leq i \leq n\}$ is admissible iff there is a $w \in a_1 \cup \cdots \cup a_n$ such that $w$ complies with every $b_i \,\|\, a_i \in \mathbf{p}$.

22 Like so:

(11) 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 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.)

23 The construction was first used (again, in a quantitative setting) by Adams (1975). In the qualitative setting (well, rank-theoretic setting (Spohn 1988)) the procedure is part of system $z$ (Pearl 1990); we'll see just how in Section 8.

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

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

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

And since $w \in a_1 \cup \cdots \cup a_n$ 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 $b_i \,\|\, a_i$ and is therefore admissible. □

The promised result: $\mathbf{p}$ is consistent iff every non-empty subset of $\mathbf{p}$ is admissible. 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 3.** If $\mathbf{p} = \{b_i \,\|\, a_i \colon 1 \leq i \leq n\}$ is consistent then $\mathbf{p}$ is admissible.

*Proof.* Suppose (for reductio) that $\mathbf{p} = \{b_i \,\|\, a_i \colon 1 \leq i \leq n\}$ is consistent but not admissible. Let $\prec$ be any ordering faithful to $\mathbf{p}$. Let $x = a_1 \cup \cdots \cup a_n$. Since $\mathbf{p}$ isn't admissible, for every $w \in x$ there is a $b_i \,\|\, a_i$ 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 $b_i \,\|\, a_i$ that $w$ flouts: that is, $w \in a_i \cap \neg b_i$ (Proposition 8). Since $\prec$ is faithful to $\mathbf{p}$ and hence to $b_i \,\|\, a_i$ 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 $v \nprec w$, completing the reductio. □

**Theorem 4.** If $\mathbf{p} = \{b_i \,\|\, a_i \colon 1 \leq i \leq n\}$ is admissible then $\mathbf{p}$ is consistent.

The proof makes use of two ideas. The first is a constructed ordered partition $\pi$ of $\mathbf{p}$ and the second is a ranking of worlds that reflects the priorities encoded in the partition.

**Definition 17** (Ranking functions)**.** A function $\kappa$ is a ranking function iff:

i. $\kappa \colon W \to \mathbb{Z}_{\geq 0}$; and

ii. $\kappa^{-1}(0) \neq \emptyset$.

By extension: (i) $\kappa(a) = \min\{\kappa(w)\colon w \in a\}$; (ii) $\kappa$ is faithful to $b \parallel a$ iff $\kappa(a \cap b) < \kappa(a \cap \neg b)$ and to a set **p** iff it is faithful to each member of **p**; and (iii) $\text{best}_\kappa(a) = \{w \in a\colon \kappa(w) \leq \kappa(v) \text{ for any } v \in a\}$.[24]

*Proof of Theorem 4.* Suppose every non-empty $\mathbf{p}' \subseteq \mathbf{p}$ is admissible. Construct a partition $\pi = \langle \mathbf{p_0}, \ldots, \mathbf{p_n} \rangle$ of **p** inductively as follows:

i. $\mathbf{p_0} = \{b \parallel a\colon \mathbf{p} \text{ admits } b \parallel a\}$

ii. $\mathbf{p_{k+1}} = \left\{b \parallel a \in \mathbf{p} \setminus \bigcup_{i=0}^{k} \mathbf{p_i}\colon \mathbf{p} \setminus \bigcup_{i=0}^{k} \mathbf{p_i} \text{ admits } b \parallel a\right\}$

Define the following ranking function based on $\pi$:

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

It follows 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 **p** that $w$ flouts. Finally: $w \preceq_\pi v$ iff $\kappa_\pi(w) \leq \kappa_\pi(v)$.

To complete the proof we show that for every linearization $\prec_\pi$ of $\preceq_\pi$ and every $i$: if $b \parallel a \in \mathbf{p_i}$ then $\text{best}_{\prec_\pi}(a) \subseteq b$. 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. $\square$

The partition $\pi$ 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?

Quite: if $a \cap b \cap c \neq \emptyset$ then $\mathbf{p} = \{b \parallel a, \neg b \parallel a \cap c\}$ has this ordered partition: $\pi = \langle \mathbf{p_0}, \mathbf{p_1} \rangle$ where $\mathbf{p_0} = \{b \parallel a\}$ and $\mathbf{p_1} = \{\neg b \parallel a \cap c\}$. Similarly, given a consistent **p** that contains $b \parallel a$ it is all too easy to find a $c$ such that $\mathbf{p} \cup \{\neg b \parallel a \cap c\}$ is consistent, too.

Analysis 1 goes wrong by consulting orderings outside those faithful to a predicament. Analysis 2 gets things right by insisting that only faithful global orderings count. It also insists that all faithful orderings matter, but

24 As with orderings, so too with ranking functions: $\kappa$ is faithful to $b \parallel a$ iff $\text{best}_\kappa(a) \subseteq b$.

this turns out to be why the analysis goes wrong. From the perspective of the template, we are running low on degrees of freedom. This is definitely suboptimal.[25]

## 8 Constraining the constraints

If the problem is that faithfulness allows too many global orderings through the front door, then perhaps we should constrain things further. This is not easy. That is because we have been assuming throughout that the set **S** of predicaments doesn't dictate what sorts of normative considerations you might face: since there are no a priori constraints on what local preferences you might have, there are likewise no a priori constraints on the space of global orderings.[26] So an analysis that goes beyond Analysis 2 has to put some boundaries in place.

The partition used in the proof of Theorem 4 is a way of enforcing boundaries by, in effect, ranking local preferences. If $b \parallel a \in \mathbf{p_i}$ and $b' \parallel a' \in \mathbf{p_j}$ where $j > i$ then $b' \parallel a'$ carries more weight than $b \parallel a$ does. You can see this in $\kappa_\pi$ and therefore in the constructed ordering: it penalizes worlds for flouting local preferences with more weight. Perhaps $\Box b$ should quantify over only the best worlds compatible with what you know in the linearizations of such constrained rankings.[27]

**Analysis 3.** Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ be any state. And $\preceq_\pi$ its induced ordering based on the ordered partition $\pi$ of $\mathbf{p}$.

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

for every linearization $\prec$ of $\preceq_\pi$.

This agrees with Analysis 2 as far as *ought* is concerned: whenever $\Box b$ is true in **s** according to Analysis 2 then it is true in **s** according to Analysis 3.

25 It also means that what might go for mushy credences 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).

26 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).

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

**Proposition 9.** For any state $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ and $b$: if $s \models_2 \Box b$ then $s \models_3 \Box b$.

*Proof.* Consider any $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ such that $s \models_2 \Box b$. So: $\text{best}_{\prec}(\cap \mathbf{k}) \subseteq b$ for every $\prec$ faithful to $\mathbf{p}$. Consider any $\prec'$ that linearizes $\preceq_\pi$: $\prec'$ is faithful to $\mathbf{p}$ (Theorem 4) and so $\text{best}_{\prec'}(\cap \mathbf{k}) \subseteq b$ and hence $\mathbf{s} \models_3 \Box b$. □

This relationship is asymmetric when it comes to *ought*s (as opposed to their negations). The place to see this is in a concrete example like (10).

**Proposition 10.** Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a, c\}$ and $\mathbf{p} = \{b \,\|\, \top, \neg b \,\|\, a\}$. Moreover, let $\mathbf{s}' = \langle \mathbf{k}, \mathbf{p}' \rangle$ where $\mathbf{p}' = \mathbf{p} \cup \{b \,\|\, a \cap c\}$. Then:

i. $\mathbf{s} \models_3 \Box \neg b$; and

ii. $\mathbf{s}' \models_3 \Box b$.

*Proof.* The ordered partition $\pi$ of $\mathbf{p}$ is $\langle \mathbf{p_0}, \mathbf{p_1} \rangle$ where $\mathbf{p_0} = \{b \,\|\, \top\}$ and $\mathbf{p_1} = \{\neg b \,\|\, a\}$. Consider any $v \in (a \cap c)$ such that $v \in b$. We will see that any $w \in a \cap c$ such that $w \in \neg b$ dominates it (with respect $\preceq_\pi$). Since $v \in a \cap b$ it flouts $\neg b \,\|\, a$ and hence $\kappa_\pi(v) = 2$. Since $w \in a \cap \neg b$ it complies with $\neg b \,\|\, a$ and flouts $b \,\|\, \top$ and hence $\kappa_\pi(w) = 1$. So in every linearization $\prec$ of $\preceq_\pi$: $w \prec v$. And so $\text{best}_{\prec}(a \cap c) \subseteq \neg b$.

Similarly, the partition $\pi'$ of $\mathbf{p}'$ is $\langle \pi, \mathbf{p_2} \rangle$ where $\mathbf{p_2} = \{b \,\|\, a \cap c\}$. Note that if $w \in (a \cap c) \cap \neg b$ then $\kappa_\pi(w) = 3$: $w$ flouts $b \,\|\, a \cap c$. And if $v \in (a \cap c) \cap b$ then $\kappa_\pi(v) = 2$: $v$ complies with $b \,\|\, a \cap c$ and flouts $\neg b \,\|\, a$. Hence in every linearization $\prec$ of $\preceq_{\pi'}$: $v \prec w$ and so $\text{best}_{\prec}(a \cap c) \subseteq b$. □

**Corollary 5.** $\Box b$ isn't persistent in either $\mathbf{p}$ or $\mathbf{k}$ with respect to $\models_3$.

*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 $\mathbf{p} = \{b \,\|\, \top, \neg b \,\|\, a\}$. Non-persistence in $\mathbf{p}$: see Proposition 10. □

This might seem like progress but, I think, the seeming is where it stops. By the letter of the law, Analysis 3 fills in the template but in spirit we have moved a long way from it. Even so, the analysis on offer for *ought*s is both too strong and too weak. The takeaway is that the problem of linking (expectation) *ought* to preference is hard.

The analysis 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 what constraints carry more weight. It does this by ignoring some global preference orderings that are compatible with the local preferences being modeled. Among the ignored orderings are those that are faithful to merely possible constraints. It does this by minimizing the suboptimality represented.[28] 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.

More precisely: in the induced $\kappa_\pi$, and so the orderings they determine, every world occupies the best position it possibly can.

**Definition 18** (Improvements)**.** Let $\kappa$ be any ranking function and $w$ any world. $\kappa'$ is a <u>(solo) improvement</u> over $\kappa$ 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 $\kappa$: $\kappa'$ makes exactly one world better without disturbing the ranking of any other world.

**Lemma.** Let $\pi = \langle \mathbf{p_0}, \ldots, \mathbf{p_n} \rangle$ be the ordered partition of $\mathbf{p}$. If $b \,\|\, a \in \mathbf{p_i}$ then $\kappa_\pi(a \cap b) = i$.

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

28 This basic idea is implemented in different ways in 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 6 was first proved in Pearl 1990.

**Theorem 6.** Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ be any state, $\pi$ the ordered partition of $\mathbf{p}$, and $\kappa_\pi$ the associated ranking function. If $\kappa$ is an improvement over $\kappa_\pi$ then $\kappa$ is not faithful to $\mathbf{p}$.

*Proof.* Suppose $\kappa$ is an improvement over $\kappa_\pi$. So for some $w$: $\kappa_\pi(w) = i + 1$ 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 the lemma that $\kappa_\pi(a \cap b) = i$. Finally, since $w \notin \text{best}_{\kappa_\pi}(a \cap b)$, we have that $\kappa(a \cap b) \nless \kappa(a \cap \neg b)$ and so $\kappa$ isn't faithful to $\mathbf{p}$. $\square$

In terms of $\preceq_\pi$ that $\kappa_\pi$ determines: modulo the local preferences, every world is assumed to be as awesome as it can be. Alas the world as we know it does not justify this cockeyed optimism. Would that it were different.

Analysis 3 is also too weak, again in two ways. First: because the way it enforces priority for more specific local preferences is too rigid. In situations like (8) it seems like more specific constraints (*Better not to go to the pub given it's Sunday*) should trump less specific ones (*Better to go to the pub than not*). The ordered partition enforces this by prioritizing the more specific constraint and so worlds that flout it are worse ceteris paribus than worlds that comply with it. But as we saw in (9) it can happen that what is more specific than what is wholly contingent, depending essentially on what you know. Here is the example again:

(9) Better to go to the pub given Alex is going. $b \parallel a$
Better to not go to the pub given Chris is going. $\neg b \parallel c$
Chris goes. $c$
Either Chris doesn't go or Alex goes. $\neg c \cup a$
You ought not go to the pub. $\mathbf{s} \models \Box \neg b$

The underlying problem shows up in two ways. Both local preferences get lumped together in the same (trivial) partition cell. This combines with what you know to over-generate competition and conflict between the local preferences. You can see this in the ordering $\preceq_\pi$ generated by this ranking (Figure 8): $\preceq_\pi$ carves the possibilities by their rank, equivalence-class mates are joined by dashed arrows, solid arrows are strict preferences between equivalence classes. The worlds compatible with $\{\neg c \cup a, c\}$ are $w$ and $y$. They are belong to the same equivalence class, but one is a $b$ world and one isn't. Hence Analysis 3 doesn't recognize this as a case of over-riding. Officially:

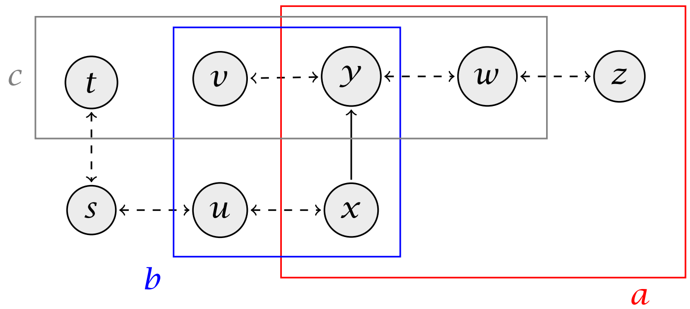


**Figure 8** Contingent specificity: induced $\preceq_\pi$

**Proposition 11.** Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{c, \neg c \cup a\}$ and $\mathbf{p} = \{b \,\|\, a, \neg b \,\|\, c\}$. Then $\mathbf{s} \not\models_3 \Box\neg b$.

*Proof.* The set $\mathbf{p}$ admits both $b \,\|\, a$ and $\neg b \,\|\, c$ and so the ordered partition is $\pi = \langle \mathbf{p_0} \rangle$ where $\mathbf{p_0} = \mathbf{p}$. Given this partition, $\kappa_\pi^{-1}(0) = \{t, s, u, x\}$ and $\kappa_\pi^{-1}(1) = \{v, y, w, z\}$. This gives the (pretty uninteresting) ordering in Figure 8. Since $\cap\mathbf{k} = \{w, y\}$ it follows that $\mathbf{s} \not\models_3 \Box\neg b$ and $\mathbf{s} \not\models_3 \Box b$. □

The second way in which Analysis 3 is too weak: it adopts a contagion view of over-riding local preferences, thereby under-generating obligations. This happens when we have multiple local preferences with the same triggering information and one of them is over-ridden.

An example, with a twist: 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.[29]

29 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:

| | | |
|---|---|---|
| (12) | It's better to wear sunglasses, given a run. | $b \,\|\, a$ |
| | It's better to go early. | $c \,\|\, a$ |
| | Except in January: it's better to go later. | $\neg c \,\|\, a \cap d$ |
| | You're going out for a run in January. | $a \cap d$ |
| | You ought to wear sunglasses. | $\mathbf{s} \models \Box b$ |

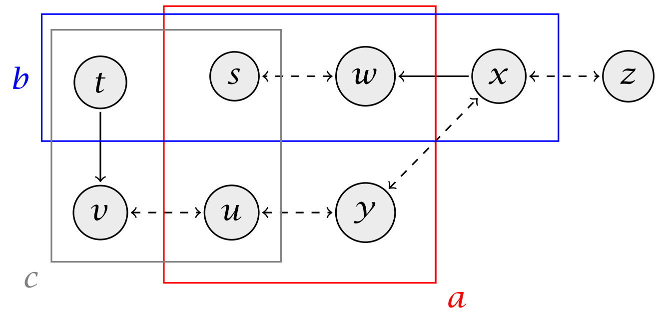


**Figure 9** Suboptimal subclass inheritance: induced $\preceq_\pi$

(13) 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. $\mathbf{s} \models \Box c$

The twist: as before, what counts as over-riding or more specific can be a contingent thing. (The same is true for when two local preferences have the same trigger; but you get the point.) Suppose, as a matter of brute fact, that candidates only declare for the race on Sundays.

(14) It's better to go to the pub, given that you're campaigning. $b \parallel d$
It's better to chat up the locals, given that you're campaigning. $c \parallel d$
It's better to not go to the pub, given it's Sunday. $\neg b \parallel a$
Either it's not Sunday or you are campaigning. $\neg a \cup d$
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. $\mathbf{s} \models \Box c$

The difference between the example and the twist should be in the noise: a good analysis should treat them the same. And in both cases: you ought to chat up the locals and you ought to stay away from the pub. Analysis 3 treats the cases differently (not good) and doesn't predict the right *ought*s (really not good).

Take (13) first. Neither of the unconditional local preferences $b \parallel \top$ and $c \parallel \top$ over-rides the other, and so they are cellmates in $\mathbf{p_0}$. But $\neg b \parallel a$ over-rides $b \parallel \top$: this gets reflected by it occupying $\mathbf{p_1}$. 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 $b \parallel \top$. 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 goodness of $c$.

In terms of the induced ordering $\preceq_\pi$: $t$ is the lone best-simpliciter world, with $\{v, u, y, x, z\}$ the rank 1 equivalence class, and $\{w, s\}$ the rank 2 equivalence class. (See Figure 9, where dashed arrows join equivalence class mates and solid arrows are strict preferences between (members of) classes.) The relevant thing: within $a$, the best worlds are $u$ and $y$ and while they are both $\neg b$-worlds, only one is a $c$-world.

**Proposition 12.** 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 $\mathbf{s} \models_3 \Box \neg b$ but $\mathbf{s} \not\models_3 \Box c$.

*Proof.* The set $\mathbf{p}$ admits both $b \parallel \top$ and $b \parallel \top$ but not $\neg b \parallel a$ 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 $\mathbf{p_1} = \{\neg b \parallel a\}$. Given this partition, $\kappa_\pi^{-1}(0) = \{t\}$ and $\kappa_\pi^{-1}(1) = \{u, v, x, y, z\}$ and $\kappa_\pi^{-1}(2) = \{s, w\}$. This gives the (again, pretty uninteresting) ordering in Figure 9. Since $\cap \mathbf{k} = \{s, u, w, y\}$ it follows that $\mathbf{s} \models_3 \Box \neg b$ and $\mathbf{s} \not\models_3 \Box c$. □

Analysis 3 doesn't cope with contingent specificity: it lumps over-riding local preferences together as competing preferences. Combine this with independence between local preferences, as in (14), and the results are not good. When what is more specific depends on what you know, Analysis 3 mistakenly lumps local preferences into the same partition cell. Focusing in on just the worlds in $d$, this generates an almost completely trivial ordering: one world is in $\kappa_\pi^{-1}(0)$ and all others are in $\kappa_\pi^{-1}(1)$. To see it graphically: take Figure 9; this represents just the worlds in $d$; finally, change the strict preference of $x$ over $w$ to the dashed double-arrow marking equivalence classhood. Zooming in further to the $a \cap d$-worlds eliminates the sole world in $\kappa_\pi^{-1}(0)$, and leaves us with a completely heterogenous set of best worlds and no clear verdicts about what you ought to do. Officially:

**Proposition 13.** Let $\mathbf{s} = \langle \mathbf{k}, \mathbf{p} \rangle$ where $\mathbf{k} = \{a, \neg a \cup d\}$ and $\mathbf{p} = \{b \parallel d, c \parallel d, \neg b \parallel a\}$. Then $\mathbf{s} \not\models_3 \Box b$, $\mathbf{s} \not\models_3 \Box \neg b$, $\mathbf{s} \not\models_3 \Box c$, and $\mathbf{s} \not\models_3 \Box \neg c$.

*Proof.* The set $\mathbf{p}$ admits all of the local preferences and therefore the ordered partition is $\pi = \langle \mathbf{p_o} \rangle$ where $\mathbf{p_o} = \mathbf{p}$. Given this partition, and taking the restriction to $d$: $\kappa_\pi^{-1}(0) = \{t\}$ and $\kappa_\pi^{-1}(1) = \{s, u, v, w, x, y, z\}$. By changing the strict preference in Figure 9 between $x$ and $w$ to equivalence, the graph models this ordering with $d$. Since $\cap\mathbf{k} = \{s, u, w, y\}$ it follows that none of $\Box b, \Box\neg b, \Box c, \Box\neg c$ are forced in $\mathbf{s}$ by $\models_3$. □

## 9 State of play

We began with a template, an observation, and a problem. The template says that *ought* is tightly tied to a relation of what-is-better-than-what. The observation is that normal predicaments under-determine such relations but don't seem to thereby under-determine the truth of *ought*s. Coping with this is the problem. But we have a partial map of the terrain for would-be solutions highlighting features that mark what can't work and why. This is progress.

## 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.

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.

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*.

Gillies, Anthony S. 2017. Conditionals. In Alex Miller Bob Hale & Crispin Wright (eds.), *A companion to the philosophy of language*. Wiley Blackwell.

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.

Kratzer, Angelika. 1981. 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. 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 artifical 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.

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.

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.

Veltman, Frank. 1985. *Logics for conditionals*. Ph.D. thesis, University of Amsterdam.