Mind & Language: Question Sensitivity

Matt Mandelkern, NYU, February 3, 2026

---

1 Belief: Background

Hintikkan belief: $[\![S \text{ believes } p]\!]^{c,w} = 1$ *iff* $\forall w' \in B_{S,w} : [\![p]\!]^{c,w'} = 1$

$B_{\cdot,\cdot}$ maps S, w to the set of worlds compatible with S's beliefs in w.

Lockean belief: $[\![S \text{ believes } p]\!]^{c,w} = 1$ *iff* $Cr_{S,w}([\![p]\!]^c) \geq t$

where $Cr_{\cdot,\cdot}$ maps S, w to a credence function representing S's credences in w, and t is a threshold in $[0, 1]$ (usually in $(.5, 1]$)

2 Background and foreground beliefs

Yalcin (2018) belief: $[\![S \text{ believes } p]\!]^{c,w} = 1$ *iff* $\forall w' \in \bigcup B_{S,w,Q_c} : [\![p]\!]^{c,w'} = 1$

$B_{S,w,Q_c} \subseteq Q_c$, a contextually salient question. $B_{S,w,\cdot}$ may be *partial*:

> The model should include information about the richness, or lack thereof, of the possible alternatives that an agent's state of belief distinguishes, about what questions these alternatives speak to and fail to speak to... A state of belief still determines a set of (maximally specific) possible worlds, but only insofar as it determines a coarser set of possibilities drawn from the partition of logical space in question... the propositions which constitute the available information of a body of belief content are those beliefs which are foregrounded by the resolution. Background beliefs correspond to the "unavailable" information implicit in the foreground beliefs

Taken for granted here is a model of questions as partitions of logical space. Other models: partitions of a subset of logical space; sets of propositions.

In this model we have closure under entailment but not *preservation of visibility*. When $p \vDash r$, S believes r wrt Q if S believes p wrt Q, but r need not be visible wrt Q even if p is.

p is visible wrt Q when $\exists q \subseteq Q : p = \bigcup q$. Aka p is a *partial answer* to Q.

(1) The absent-minded detective believes the butler did it, but totally overlooks the possibility it was the chauffeur. Were he to consider this possibility, it might shake up his view. Does the detective believe that the chauffeur did not do it?

Stalnaker 1984

Yalcin's model: his state of mind is decided on the question {*butler, not butler*} but not on the question {*butler, chauffeur, neither*}

> Deductive thought... is generally a matter of dialing the resolution of one's state of mind into one that renders the premises and conclusion visible

Yalcin's framework is extremely flexible, but he doesn't motivate much

of that flexibility. His motivations all have to do with awareness; as far as that goes, a better model might be this:

Sidebar on presupposition and trivalence.

$$[\![S \text{ believes } p]\!]^{c,w} \text{ is defined only if } A_{S,w}(Q_c) = 1$$
$$\text{where defined, } = 1 \text{ iff } \forall w' \in B_{S,w} : [\![p]\!]^{c,w'} = 1$$

where $A_{S,w}(Q) = 1$ iff S is aware of Q in w.

Would this be enough? Suppose William believes he will not fight a war with France.

(2) William believes he will not fight a nuclear war with France.

Hoek gives more motivations for a Yalcin-style approach which also tries to incorporate cross-question constraints.

Some intuition this is false, not undefined. For Yalcin to model that, William's belief state on {*war, no war*} entails *no war* but on {*nuclear war, no nuclear war*} is defined but compatible with nuclear war.

The model does not help with believing necessary truths, but it can help distinguish beliefs in necessarily equivalent truths: even when p and q are equivalent, they may make different questions salient.

3 More question sensitive belief

(3) *Susie has 450 tickets in a fair 1000 ticket lottery. Billie has the other 550 tickets.*

a. -Do you think Susie will win? -No.

b. -Who, if anyone, do you think will win the lottery? -Billie.

(4) *Susie has 450 tickets in a fair 1000 ticket lottery. The other 550 tickets are evenly distributed among 550 other people, conveniently called 1, 2, 3, . . . , 550.*

a. -Do you think Susie will win? -No.

b. -Who, if anyone, do you think will win the lottery? -Susie.

Levi 1967; Kahneman and Tversky 1982; Windschitl and Wells 1998; Yalcin 2018; Hawthorne et al. 2016; Holguin 2020. It seems pretty safe to move back and forth between 'think' and 'believe' as in these examples; of course there are differences between the two, but as far as our interests go they seem interchangeable. A significant question here about whether this just reflects a "mistake" about probabilities.

It looks like, relative to the binary question {*Susie wins, Susie loses*}, you should think Susie will lose; relative to the 551-ary question {*Susie wins, 1 wins, 2, wins, . . . 550 wins*}, you can think Susie will win.

The contrast between (3) and (4) is also surprising and points towards question-sensitivity as well. If we take these judgments seriously:

- You can rationally believe p when your credence in p is $< .5$
- You can rationally believe p relative to Q and believe $\neg p$ relative to Q' when $Q \neq Q'$

The first point is bad for Lockeanism, even question-sensitive Lockeanism; since if $t < .5$, you can believe p and believe $\neg p$.

4 Questions and lotteries

Question-sensitivity may help explain ambivalence about closure. On the one hand, it's strange to say things like:

(5) #I think A will be there, and I think B will there, but I don't think they both will be there.

And this report seems totally merited:

(6) S: I think A will be there, and I think B will be there.
R: S thinks that A and B will be there.

This is not just an error; compare:

(7) S: There's a good chance A will be there and a good chance B will be there.
R: $\not\rightsquigarrow$S thinks there's a good chance that A and B will both be there.

On the other hand, in lottery style cases, it seems fine to believe of each ticket that it will lose, but not to put those beliefs together:

(8) I think ticket 1 will lose.
I think ticket 2 will lose.
$\vdots$
I think ticket 1,000,000 will lose.
$\not\rightsquigarrow$ I think every ticket will lose.

Possibly: for each n, relative to the question {*n will win, n will not win*}, you can believe *n will not win*; but relative to the question {*some ticket wins, no ticket wins*}, you must believe *some ticket wins*.

5 Cogency/Filteredness

This is all easy to model in Yalcin's extremely flexible framework. Holguín adds one further constraint:

Cogency: $\forall p, r \in Q_c : Cr_{S,w}(p) > Cr_{S,w}(r)$ *implies* $r \in B_{S,w,Q_c} \rightarrow p \in B_{S,w,Q_c}$

Dorst and I called this same constraint 'Filteredness' since 'Cogency' seems so generic. It is a constraint on *rational* belief.

Cogency is motivated by observations like this. Suppose you have credences $\langle .4, .3, .15, .1, .05\rangle$ in A, B, C, D, E winning. It seems any of the following are rationally permissible beliefs:

(9) a. A wins.
b. A or B wins.
c. A or B or C wins.

d. A or B or C or D wins.
e. A or B or C or D or E wins.

By contrast, it seems bizarre to believe:

(10) a. B or C wins.
b. B wins.
c. B or C wins.
d. B or C or D wins.
e. A or C wins.
⋮

etc. In other words, although rational question-sensitive belief seems relatively unconstrained, it seems like it is "constructed" this way:

For some threshold t_c, $B_{S,w,Q} = \{r \in Q : Cr_{S,w}(r) > t_c\}$.

Holguín's view diverges from this only in cases of tie-breaking. Although a rather abstruse case, tie-breaking is crucial for the question of whether belief supervenes on credence plus a contextual threshold, which is compatible with no-tie-breaking but not otherwise. Cf. Levi 1967; Shoham 1987; Lin and Kelly 2012a,b; structures with this feature are relatively common in the literature.

6 Want: Background

Hintikkan want: $[\![S \text{ wants } p]\!]^{c,w} = 1$ *iff* $\forall w' \in D_{S,w} : [\![p]\!]^{c,w'} = 1$

$D_{\cdot,\cdot}$ maps S, w to the set of worlds compatible with S's desires in w.

(11) a. I want to get better. $\not\rightarrow$ I want to have been sick.
b. I want Trump to have an unproductive second term. $\not\rightarrow$ I want Trump to have a second term.

Condoravdi etc: 'S wants p' is only defined when S doesn't believe p. Not so obvious this works:

More plausible: 'S wants p' is only defined when S doesn't believe $\overline{p}$.

(12) *Getting served a big piece of lasagna.*
This is what I want.

A different kind of case in the same ballpark:

Stalnaker, Heim

(13) a. Nicholas wants a free trip on the Concorde.
b. Nicholas wants a trip on the Concorde.

Conditional want: $[\![S \text{ wants } p]\!]^{c,w} = 1$ *iff* $\forall w' \in B_{S,w} : f(p, w') \succ f(\overline{p}, w')$

where $\succ$ is an ordering on worlds by desirability and $f(p, w)$ is a selection function representing something like relative closeness.

7 Threshold (Lockean) views

Not actually clear to me how this helps with the pair in (13). A different diagnosis: the expected value of riding on a Concorde is low, but the expected value of a free ride on a Concorde is high.

Simple Satisficing: S wants p iff the expected value S assigns to p exceeds a given threshold.

Simple Maximizing: S wants p iff the expected value S assigns to p is maximal, among a set of alternatives.

(14) *You are deciding where to go to dinner. There are three options: the pizzeria, the ramen shop, and the hot dog stand. Pizza and ramen both sound good; hot dogs don't; pizza is most appealing.*

Phillips-Brown 2021

a. I want to go to the ramen shop.
b. But I want to go to the pizzeria even more.

This speaks in favor of a Satisficing view.

(15) *Johnnie must do one of two chores: mow the lawn or wash the dishes, both of which he hates. He prefers mowing the lawn to washing the dishes.*

A: I see that Johnnie is going to mow the lawn. Is that something he wants to do?
B: No, Johnnie doesn't want to mow the lawn (he's only going to because he has to do a chore and mowing the lawn is better than washing the dishes).
C: Why is Johnnie going to the shed?
D: Johnnie wants to mow the lawn (and that's where the mower is—of course he hates mowing the lawn, but it's better than washing the dishes).

This could motivate a contextualist version of the Satisficing view where the threshold varies by context. Phillips-Brown argues for more:

Prospect Sensitivity: *S wants p* is true in c iff there is a p-prospect p_i in the prospect set in c such that: the expected value S assigns to p_i meets the threshold in c.

a p-prospect is a way for p it be true. A prospect set is thus in some sense a question, but not (in P-B's formulation at least) in the sense of partitioning logical space.

(16) *Mo likes winning, but only winning fairly. Mo has registered for an upcoming race, but he knows that faster runners than him have also registered; he can win only if he cheats.*

A: I heard Mo's not going to train very hard for this race. Doesn't he want to win?
B: Of course Mo wants to win this race. (But he can only win by cheating, which is why he's not going to train very hard.)
C. I heard Mo's not going to train very hard for this race. Doesn't he want to win?
D. No, he doesn't want to win this race. (He can only win by cheating, which is why he's not going to train very hard.)

(17) *Li is the Secretary of State, and he prefers diplomacy to violence. The*

'Different ways of Mo winning are in some sense being talked about in the different contexts.'

country faces two enemies: the pretty bad guys and the very bad guys. Li had proposed to the President three options, which in descending order of Li's preference are these: negotiate with the very bad guys, bomb the very bad guys, bomb the pretty bad guys. (Best to bomb no one, but if someone is to be bombed, it should be the very bad guys, not the pretty bad guys.) The President immediately dismisses the idea of negotiating with the very bad guys and says she'll soon decide between the remaining two options.

a. Li wants the President to bomb the VERY BAD GUYS.
b. Li wants the President to BOMB the very bad guys.

A different kind of case, from Phillips-Brown 2025:

(18) *John used to love eating lobster, but has a recently diagnosed shellfish allergy.*

a. I don't want to eat this lobster.
b. I want to eat this lobster.

This looks hard to account for in a framework like the present one, and P-B uses it to motivate a different approach on which S wants p iff p is entailed by one of the best answers to a relevant question compatible with your beliefs.

Here, the relevant question can either be {*eat a lobster, don't eat a lobster*} or {*eat a lobster and die, don't eat a lobster and die, eat a lobster and don't die, don't eat a lobster and die*}.

Prima facie, it looks like it has nothing to do with decision theory; the expected value of eating a lobster is extremely low.

But as P-B (p.c.) has pointed out, decision theory still seems to play a role in cases like this. Whether (18-b) is true may depend on your credence that the thing before you is not just a lobster but a non-rancid lobster, and so on.

References

Hawthorne, J., Rothschild, D., and Spectre, L. (2016). Belief is weak. *Philosophical Studies*, 173:1393–1404.
Holguin, B. (2020). Thinking, guessing, and believing. Manuscript, NYU.
Kahneman, D. and Tversky, A. (1982). Variants of uncertainty. *Cognition*, 11:143–157.
Levi, I. (1967). *Gambling with truth: An essay on induction and the aims of science*. MIT Press, Cambridge, MA.
Lin, H. and Kelly, K. T. (2012a). A geo-logial solution to the lottery paradox. *Synthese*, 186:531–575.
Lin, H. and Kelly, K. T. (2012b). Propositional reasoning that tracks probabilistic reasoning. *Journal of Philosophical Logic*, 41:957–981.
Phillips-Brown, M. (2025). Some-things-considered desire. *Semantics and Pragmatics*, 18(11).
Shoham, Y. (1987). A semantical approach to nonmonotonic logics. In Ginsberg, M., editor, *Readings in nonmonotonic reasoning*. Morgan Kauffman, Los Altos, CA.
Stalnaker, R. (1984). *Inquiry*. MIT.
Windschitl, P. and Wells, G. (1998). The alternative-outcomes effect. *Journal of Personality and Social Psychology*, 75(6):1411–1423.
Yalcin, S. (2018). Belief as question-sensitive. *Philosophy and Phenomenological Research*, XCVII(1):23–47.