A Puzzle About Weak Belief

1 Introduction

I present an intractable puzzle for the currently popular view that belief is weak — the view that expressions like '*S* believes *p*' ascribe to *S* a doxastic attitude towards *p* that is rationally compatible with low credence that *p* — defended in Hawthorne et al. 2016, Dorst 2019, Rothschild 2020, Dorst and Mandelkern 2021, Mandelkern and Dorst 2022, Holguín 2022 and Goodman and Holguín 2022. This view has not escaped criticism — see Moss 2019, Williamson forthcoming, Nagel 2021 and Clarke forthcoming — but the critique offered here is unique, focusing on issues concerning beliefs in conditionals. I show that proponents of weak belief either cannot consistently apply their preferred methodology when accommodating beliefs in conditionals, or they must deny that beliefs in conditionals can be used in reasoning.

§2 reviews data motivating, and then outlines, the recent promising theory of weak belief defended by Holguín (2022). §3 argues Holguín's theory generates implausible predictions concerning conditionals, and that natural replies to this problem prevent rational agents from using the conditionals they believe in reasoning. §4 outlines the puzzle in general terms, that any theory of weak belief must answer.

2 Holguín's Theory

The thesis that belief is weak is usually motivated via felicity judgements concerning various belief ascriptions. Hawthorne et al. (2016, 1395) observe that it sounds perfectly reasonable to express a belief in *p* whilst denying one has a strong doxastic attitude towards *p*:

(1) $\checkmark$ I $\left\{\begin{matrix}\text{believe}\\\text{think}\end{matrix}\right\}$ it will rain, but I $\left\{\begin{matrix}\text{don't know}\\\text{am not sure}\end{matrix}\right\}$ it will rain.[1]

Moreover, if one only has a low credence that *p*, expressions of belief can nevertheless sound unobjectionable (Hawthorne et al. 2016, 1400). If Horse *A* has a 60% chance of winning, but *B* and *C* have 20%, (2) sounds fine:

[1] I follow proponents of weak belief in treating 'thinks' and 'believes' synonymously. Whilst I do not object to this move here, others have, such as Nagel (2021).

(2) ✓ I $\begin{Bmatrix}\text{believe}\\ \text{think}\end{Bmatrix}$ *A* will win.

(2) can sound fine even if *A* is less than 50% likely to win; for example, if *A* has 40%, and *B* and *C* only 30%. According to Holguín (2022, 5-6), (2) sounds fine even when *A*'s chance winning is arbitrarily low, so long as *A* remains the favourite.

Proponents of the weak belief thesis argue that the felicity of (1) and (2) is strong evidence that believing *p* is rationally compatible with low credence that *p*.[2] I will not object to this style of argument here — see (Holguín 2022, 3) for a more thorough defence of it.[3] However, it's important to note that, if we do endorse this style of argument, further data concerning the felicity of belief acsriptions highlights additional features of rational beliefs our preferred theory should capture.

As Holguín (2022, 8) observes, it sounds bizarre to express beliefs that fail to be closed under deduction:

(3) # I believe both that Alice will come to the party and that Bob will, but it's not as if I believe both of them will come to the party.

(4) # I think that Alice will come to the party, and that if Alice will, so will Bob. But I hardly think Bob will come to the party.

Moreover, as Holguín (2022, 8) notes, "it is highly plausible that rational agents can come to think that a proposition is true by deducing it from other propositions they already think are true."

Should our theory therefore predict that rational beliefs are closed under deduction? It's unclear; further data supports the contrary (Holguín 2022, 9). Suppose Bond is 40% likely in London, 30% likely in Munich, and 30% likely in Berlin. If asked what *country* Bond is in, it's reasonable to respond:

(5) ✓ I $\begin{Bmatrix}\text{believe}\\ \text{think}\end{Bmatrix}$ Bond is in Germany.

If asked what *city* Bond is in, it's reasonable to respond:

(6) ✓ I $\begin{Bmatrix}\text{believe}\\ \text{think}\end{Bmatrix}$ Bond is in London.

It's nevertheless impossible to hear a belief in the conjunction as rational:

[2]See Hawthorne et al. 2016, Rothschild 2020 and Holguín 2022 for further data of this kind in support of weak belief.

[3]One complaint is that the locution "I think" is used to hedge an assertion, rather than as a self-ascription of belief (Nagel 2021, 53-4). However, the data considered here applies equally well when considering third-person ascriptions like "John thinks horse *A* will win."

(7) # I $\left\{\begin{matrix}\text{believe}\\ \text{think}\end{matrix}\right\}$ Bond is in Germany and in London.

Impressively, Holguín (2022, 12-13) proposes a theory that can accommodate all of our judgements concerning (1)-(7). Holguín claims beliefs are question-sensitive, and that one can rationally believe p, with respect to question Q, iff p follows from one's best guess to Q.[4] Let W be our set of worlds, propositions p be subsets of W, S be our agent, e (a proposition) S's total evidence, and C be S's probabilistically coherent credence function which is certain of e. A question, Q, is a partition of W. For each context of attribution we assume there is one question-under-discussion (QUD) Q which determines the referent of 'thinks' (thinks_Q) and 'believes' (believes_Q). (I'll drop the Q-parameter when there is no risk of ambiguity.) $g_{C,Q}$ is S's "best guess" to Q: the cell of Q that C assigns greatest probability to (or the union of cells tied for greatest probability). Here's Holguín's theory:

Best Guess Account (BGA)
S is rationally permitted to believe$_Q$ *p iff* $e\&g_{C,Q}$ *entails p.*

BGA satisfies the data regarding (1) and (2): one is permitted to believe p despite having low credence that p if p is jointly entailed by one's best guess and evidence. Explaining (3)-(7) is more subtle. BGA identifies a unique strongest proposition one can permissibly believe_Q: $e\&g_{C,Q}$. What one can permissibly believe is deductively closed with respect to Q, since one can believe_Q p iff p is entailed by $e\&g_{C,Q}$. That's why (3), (4), and (7) sound bad.[5] Why do (5) and (6) sound *good*? Because they sound good *only with respect to different questions*. With respect to what country Bond is in, (5) sounds good but (6) bad. With respect to what city Bond is in, (6) sounds good but (5) bad. Apparent failures of closure are merely cross-contextual failures.

So far, so good. Trouble arises when we consider beliefs in conditionals.

3 Conditionals

We often refrain from giving categorical answers to questions, instead only stating what we think given some condition. Suppose City will play United, but there's doubt concerning the fitness of City's star striker, Haaland. If asked whether United will win, I may refrain from offering a categorical opinion and simply assert:

(8) ✓ I $\left\{\begin{matrix}\text{belive}\\ \text{think}\end{matrix}\right\}$ that if Haaland starts, United will lose.

[4] Beliefs also need to be "cogent" — see (Holguín 2022, 18) and (Dorst and Mandelkern 2021, 586) — but this complication does not matter for our purposes.

[5] In fact, Holguín needs to say more here. BGA only *permits* one to believe any deductive consequences of one's beliefs. To fully explain (3), (4) and (7) Holguín must also endorse some condition like: if S permissibly *believes* $p_1, ...p_n$, and $p_1, ...p_n$ jointly entail q, then S is *required* to believe q.

These beliefs in conditionals look as weak as regular beliefs:

(9) ✓ I $\left\{\begin{array}{c}\text{don't know}\\ \text{am not sure of}\end{array}\right\}$ this, but I $\left\{\begin{array}{c}\text{belive}\\ \text{think}\end{array}\right\}$ that if Haaland starts, United will lose.

Moreover, we often use our beliefs in conditionals to reason towards further conclusions:

(10) ✓ I think that if Haaland starts, City will win. If he's injured, then I think United will at most get a draw. Either way, then, I don't think United will win.

Such reasoning is essential when answers to the QUD fail to have salient probabilities. Finally, one can felicity hedge a belief by asserting it in conjunction with a belief in a conditional whose consequent conflicts with the hedged belief:

(11) ✓ I believe United will win; nevertheless, I believe that if Haaland starts, United will lose.

Applying the methodology in §2, our theory of weak belief ought to accommodate this data. Assuming beliefs in conditionals are beliefs in propositions,[6] it's easy to apply BGA: S is permitted to believe$_Q$ a conditional iff it is entailed by $e\&g_{C,Q}$. Consider a simple example:

> Six-Sided Die. Alice and Bob will roll a fair six-sided die. Alice wins if it lands 1-4; Bob wins if it lands 5-6.

No matter the QUD, BGA makes implausible predictions here. First, suppose the QUD is $Q_{Who?}$: {*Alice will win*; *Bob will win*}; intuitively: '*Who will win*?'. Assuming S's evidence offers no relevant information beyond that specified in Six-sided Die, since Alice is $\frac{2}{3}$rds likely to win, $e\&g_{C,Q}$ is *Alice will win*, meaning BGA desirably predicts the following expressed belief is rational for S:

(12) ✓ I think Alice will win.

Now consider what conditionals S is permitted to believe with respect to $Q_{Who?}$. If the die lands greater than 3, the probabilities are reversed: Bob is now $\frac{2}{3}$rds likely to win. Accordingly, the belief expressed in (13) looks rational:

(13) ✓ I think that if the die lands greater than 3, Bob will win.

[6] A controversial assumption I can't defend here; see Edgington 1995.

Indeed, this judgement is on as solid ground as the foundational judgements in support of weak belief, (1) and (2). Yet BGA predicts (13) is rational only if S's best guess — *Alice will win* — entails *if the die lands greater than three, Bob will win*. No sensible theory of indicative conditionals allows for this entailment. The world in which the die lands 4 is a world in which Alice wins, yet *If the die lands greater than* 3*, Bob will win* has a true antecedent and false consequent, which any sensible theory of conditionals will take as sufficient for its falsity at that world. BGA therefore predicts (13) expresses an *irrational* belief.

Perhaps BGA does better when shifting the QUD to a question about a conditional, $Q_{>3?}$: {*If the die land greater than* 3*, Bob wins*; *If the die lands greater than* 3*, Alice wins* };[7] intuitively: '*Who will win if the die lands greater than* 3*?*'. Deriving what one's best guess with respect to $Q_{>3?}$ requires making controversial assumptions about the probability of conditionals.[8] To sidestep this, I'll charitably assume that $g_{C,Q_{>3?}}$ is *if the die lands greater than three, Bob will win*, delivering the result that the belief expressed in (13) is rational.

The problem is that, with respect to $Q_{>3?}$, the belief expressed by (12) is now irrational. BGA predicts S permissibly believes$_{Q_{>3?}}$ *Alice will win* iff it is entailed by *If the die lands greater than* 3*, Bob will win*. Again, no sensible theory of indicative conditionals will predict this. This entailment can only hold if the conditional is false at worlds in which the die lands 5 or 6, since those are worlds in which Alice does not win. That cannot happen on theories for which the truth of an antecedent and consequent at a world is sufficient for the truth of the corresponding conditional at that world, such as according to the theory in Stalnaker 1968. And whilst theories that demand more for the truth of a conditional can *consistently* model the required entailment, such models are bizarre.[9] BGA predicts (13) expresses a rational belief only at the cost of no longer predicting the same about (12).

I will now consider two replies to my arguments so far. The first reply attempts to fully utilise Holguín's question-sensitive framework. The second modifies BGA by proposing that rational beliefs in conditionals are determined by corresponding conditional probabilities.

---

[7]That $Q_{>3?}$ is a partition requires controversially assuming Conditional Excluded Middle, validated according to the theory in Stalnaker 1968 (cf. Stalnaker 1980) but not, for instance, the theory in Kratzer 1986. My arguments don't depend on it.

[8]A controversy stirred by the vast literature on triviality results, kick-started by Lewis (1976).

[9]Consider the theory in Kratzer 1986, whereby *If the die lands greater than* 3*, Bob wins* is true at w iff at the set of worlds determined by w, $b(w)$, all the greater-than-three-worlds are such that Bob wins. Call a world in which the die lands n an 'n-world'. For *If the die lands greater than* 3*, Bob wins* to entail *Alice Wins*, the former must be false when Bob wins. Hence, when w is a 5-world or 6-world, $b(w)$ must contain some 4-world. Now, *If the die lands greater than* 3*, Bob wins* is not a contradiction, and so must be true at some world, v. $b(v)$, however, *cannot* contain any 4-world, since otherwise the conditional will be false at v. So, we can only consistently model the required entailment if we, bizarrely, force $b(w)$ to contain 4-worlds when w is a 5-or-6-world, but forbid this for 1-to-4-worlds in which the conditional is true.

3.1 Question-sensitivity

Perhaps, when determining whether an agent is rational to believe p, we tend to asses this with respect to the QUD $\{p, \neg p\}$ (or some QUD for which p is a complete answer). If so, BGA makes the correct predictions — S permissibly believes both *If the die lands greater than 3, Bob will win* (with respect to $Q_{>3?}$) and S permissibly believes *Alice will win* (with respect to $Q_{Win?}$) — and all I've said above is a feature, not a bug, of BGA.

Two objections. First, (12) and (13) sound good not only in isolation, but also within the same breath. The effect is similar to the hedging assertion of (11):

(14) ✓ I think Alice will win; nevertheless, I think that if the die lands greater than 3, Bob will win.[10]

Yet there is no QUD where BGA predicts this conjunction of beliefs as permissible. We have already seen neither $Q_{Win?}$ nor $Q_{>3}$ allow for both beliefs expressed in (14) to be rational.[11] A further contender is the conjunctive question $Q_{Win?\&>3}$: *'Who will win, and who will win if the die lands greater than 3?'*; however, for this question it's not even obvious *either* belief counts as rational by BGA. Since the details are thorny — requiring us to make concrete assumptions about the probabilities of conditionals — I've placed them in the following footnote.[12]

[10] As a referee observes, (14) can also sound felicitous if constructed as a single belief in a conjunction, like: "I think Alice will win, but not if the die lands greater than 3".

[11] Might we, as a referee suggests, allow for mid-context shifts, as does Dorr (2014), along with Goodman and Lederman (2021)? Such a view needs to tread carefully to avoid over-generating. It's unclear how a theory that allows for mid-sentence context shifts can, in a principled manner, predict the felicity of (14), without also falsely predicting that bad-sounding ascriptions should sound fine, like (15) and (16) outlined below, or, say, the sentence "I believe Bond is in Germany and I believe Bond is in London."

[12] Determining one's best guess to $Q_{Win?\&>3}$ requires making concrete assumptions about the semantics of conditionals. I'll be setting aside the counter-intuitive material conditional approach, although see Williamson 2020 for a recent defence. Approaches like Kratzer (1986), whereby the truth of a conditional requires all relevant antecedent-worlds to be consequent-worlds, also perform badly here: so long as a 4-world is relevant, *If the die lands greater than* 3, *Bob will win* has probability 0.

What about the theory in Stalnaker 1968, whereby a conditional is true just in case the closest antecedent-world is a consequent-world? Let's suppose $Q_{Win?\&>3}$ is the partition: {*Alice Wins & If greater than* 3, *Bob wins*; *Alice Wins & If greater than* 3, *Alice wins*; *Bob Wins & If greater than* 3, *Bob wins*; *Bob Wins & If greater than* 3, *Alice wins* }. To calculate the probability of *Alice Wins & If greater than* 3, *Bob wins*, observe first that this proposition is false whenever the die lands greater than 3, meaning it has at most probability $\frac{1}{2}$. Where $f(a, w)$ is the closest a-world to w, and '> 3' the proposition *The die lands greater than three*, observe further that the proposition in question is false at any 1-3-world w in which $f(>3, w)$ is a 4-world. Making the natural assumption that $f(>3, w)$ is a 4-world for $\frac{1}{3}$rd of the 1-3 worlds, we are left with just $\frac{2}{3}$rds of the 1-3 worlds in which our proposition is true, meaning it has total probability of $\frac{1}{2} \times \frac{2}{3} = \frac{1}{3}$. A similar calculation yields that *Alice Wins & If greater than* 3, *Alice wins also* has probability $\frac{1}{3}$, meaning BGA doesn't predict that the beliefs in (14) are entailed by one's best guess, and hence are not rationally permissible.

Can Holguín reply by allowing one's best guess to break ties, as per Dorst and Mandelkern 2021? On such a view, our agent *would* be permitted to believe *Alice wins, and if the die lands greater than 3, Bob wins*, since it

Second, if assessing a belief in a conditional requires changing the QUD, we generate a conflict with our use of conditionals in reasoning, as with (10). If assessing a conditional requires changing the QUD, which is turn requires changing the semantic values of "thinks" and "believes", we cannot vindicate this practice: if the QUD is Q_1, the fact I rationally think$_{Q_2}$ that some conditional is true is just irrelevant to what I can think$_{Q1}$.

A natural reply to this objection is to posit bridge principles allowing what one can permissibly think$_{Q_1}$ to impact what one can permissibly think$_{Q_2}$. However, doing so will undermine what made positing question-sensitivity initially attractive. Recall the discussion of (5)-(7): we require a substantive independence between what one can permissibly think with respect to different questions in order to avoid predicting that the belief expressed in (7) can ever be rational. It's not obvious what independently motivated bridge principles will allow for one to rationally obtain the belief (14) from the beliefs in (12) and (13), yet prevent one from likewise obtaining the belief in (7) from the beliefs in (5) and (6).[13]

3.2 Conditional Probabilities

This reply attempts to make the correct predictions concerning which beliefs in conditionals are rational all with respect to a single QUD. Where $C(\cdot \mid \cdot)$ is S's conditional credence function ($C(q \mid p) = \frac{C(p\&q)}{C(p)}$ for any p such that $C(p) > 0$), $g_{C_p,Q}$ is the cell of Q that $C(\cdot \mid p)$ assigns highest probability to (or the union of cells with highest probability), and '$p \rightarrow q$' picks out the proposition expressed by 'If p, q' (relative to the context), consider:

Conditional Best Guess Account (CBGA)
S is rationally permitted to believe$_Q$ $p \rightarrow q$ *iff* $e\&g_{C_p,Q}$ *entails* q.

Assuming any proposition q is equivalent to $\top \rightarrow q$, $\top$ the tautology, CBGA addresses both what conditional and categorical propositions S can permissibly believe.

With respect to just $Q_{Win?}$, CBGA desirably predicts the beliefs expressed in (12)-(14) are rational. $g_{C_\top,Q_{Win?}}$ = *Alice will win*, meaning one is permitted to believe Alice will win, yet, where '> 3' stands for *the die lands greater than* 3, $g_{C_{>3},Q_{Win?}}$ = *Bob will win*, meaning one is, with respect to $Q_{Win?}$, also permitted to believe *If the die lands greater than* 3, *Bob will win.*

is an answer with tied-for-highest probability. However, this move would not help solve the problem: we can just alter the case such that the die is weighted slightly in favour of landing on 4, making *Alice wins, and if the die lands greater than 3, Alice wins* the strictly most probable answer.

[13] The bridge principles in Hoek forthcoming are of no help. Hoek suggests that if one can form beliefs about a conjunctive question, those beliefs should cohere with one's beliefs about the separate conjunct questions. Assuming one does not have the belief in (7), this means one cannot have both the beliefs in (5) and (6).

However, in predicting (12)-(14), CBGA goes out of the frying pan and into the fire: permissible beliefs now fail to be closed over entailment in various ways. Consider:

Modus Tollens
$(p \rightarrow q)$ and $\neg q$ jointly entail $\neg p$.

In Six-sided Die, whilst S can permissibly believe *Alice will Win* and *If the die lands greater than* 3*, Bob will win*, S cannot permissibly believe *The die will land greater than* 3. CBGA has this verdict since $g_{C_\top, Q_{Win?}}$ is consistent with the die landing on 4. Hence S's permissible beliefs are not closed under Modus Tollens.

More trouble. CBGA predicts S's beliefs are not closed under:

Proof By Cases
If $(p_1 \vee p_2)$ is a tautology, then $(p_1 \rightarrow q)$ and $(p_2 \rightarrow q)$ jointly entail q.

Consider:

> Ten-sided Die. Alice, Bob and Charlie are rolling a fair ten-sided die. Alice wins if it lands 1-3, Bob if it lands 4-7, Charlie if it lands 8-10.

Assuming S's credences conform to the chances, with respect to $Q_{Win?}$ — in this case, {*Alice will win*; *Bob will win*; *Charlie will win*} — CBGA tells us S is permitted to believe *If the die lands greater than* 5*, Alice will win* and *If the die does not land greater than* 5*, Charlie will win*, yet S is not permitted to believe *Bob will lose*. Indeed, CBGA predicts one can believe Bob will win!

Similar arguments apply to further principles I do not space to discuss. Not even closure under Modus Ponens survives: allowing q is itself to be a conditional, S may believe permissibly p and $p \rightarrow q$ without believing q.[14]

These are bad results, for two reasons. First, theories of weak belief are motivated, as in §2, by data concerning what belief ascriptions sound good or bad. As we saw with (3) and (4), ascriptions of beliefs that fail to be deductively closed generally sound bad. The same is true with the above failures of closure:

(15) # I think that, if the die lands greater than 3, Bob will win, and I think Bob won't win. But I wouldn't say I think the die won't land greater than 3.

(16) # I think that, if the die lands greater than 5, Charlie will win; otherwise, I think Alice will win. Nevertheless, I think Bob will take it.

[14]E.g. for Six-sided Die, let the QUD be $Q_{Win?}$, p be *Alice will win*, and q be *If the die lands greater than* 3*, Alice will win*; cf. McGee 1985.

Meanwhile, since utterances of (12)-(14) sound *good* — as do similar sentences regarding TEN-SIDED DIE like "I think that if the die lands greater than 5, Charlie will win" — it becomes impossible for the proponent of weak belief to capture all the data that was initially intended to motivate their theory, no matter what fix to BGA we offer.

Second, as with (10), it is common practice to give conditional answers to questions to aid reasoning towards a conclusion. Such reasoning will make use of principles like MODUS TOLLENS and PROOF BY CASES. The more principles of conditional logic we contravene, the harder it is to vindicate this essential practice of ours.

4 The Puzzle

My discussion focused on Holguín's theory of weak belief and amendments to it. However, the puzzle is in fact quite general.

On the one hand, just as it seems one can rationally have weak beliefs, it seems one can rationally have weak beliefs in conditionals. These judgements include, for example, that in SIX-SIDED DIE one can rationally believe *If the die lands greater than three, Bob will win*, and that in TEN-SIDED DIE one can rationally believe *If the die lands on five or less, Alice will win.* Call the desiderata for a theory to predict these weak beliefs in conditionals as rational '*Conditional Weakness*'.

On the other hand, it is a common rational practice of ours to use beliefs in conditionals to draw further conclusions. Indeed, expressions of belief that demonstrate one not to have drawn obvious inferences from propositions one already believes sound patently irrational. Call the desiderata for a theory to respect these facts '*Closure*'.

The puzzle is that we cannot have both *Conditional Weakness* and *Closure*. This is demonstrated by the above cases, SIX-SIDED DIE and TEN-SIDED DIE: as soon as we have a theory that allows an agent to rationally believe the relevant conditionals, we have a theory that violates principles of conditional logic like MODUS TOLLENS and PROOF BY CASES.[15]

[15] A referee suggests it may be useful to distinguish between beliefs being permissible *in isolation* from beliefs being *collectively* permissible alongside other beliefs. Perhaps Holguin's view can give us an account of isolated permissible belief, whereas collective permissibility requires one's beliefs to satisfy something like the constraints in Hoek forthcoming constraints (see fn. 13). So, the felicity of, say, (5) and (6) is explained as they are both rational is isolation, yet an agent may not be rational in having *both* beliefs expressed in (5) and (6), since doing so violates a Hoek-like constraint.

I doubt this move can work for conditionals like (12) and (13). The hope would be that while these beliefs are rational considered in isolation, they are not rational when considered together due to a Hoek-like constraint, blocking the prediction that (15) expresses a rational belief state. The problem is that this move will have trouble explaining why (14) nevertheless sounds felicitous. However, I think distinguishing between isolated versus collectively permissible beliefs is a promising strategy for further research.

Can we give up *Conditional Weakness*? Only if we have some pragmatic explanation why weak beliefs in conditionals sound, but in fact are not, rational. I cannot see what such an explanation could look like that would not undermine the original data supporting rational weak belief.

Can we give up *Closure*? Again, only if we have some pragmatic story why expressions of beliefs that fail to be deductively closed sound irrational, even though they needn't be. Moreover, to vindicate our common practice of using beliefs in conditionals to deduce further conclusions, we need to posit restricted closure principles that allow agents to reason with conditionals in most cases. Whilst I do not have a knock-down argument this cannot be done, I cannot see what the requisite pragmatic story or restricted closure principles would look like.[16]

Massachusetts Institute of Technology
USA
pearsonj@mit.edu

References

Clarke, Roger, forthcoming. 'Strong Belief is Ordinary'. *Episteme*, 1–21.

Dorr, Cian, 2014. 'Transparency and the Context-Sensitivity of Attitude Reports'. In Manuel García-Carpintero and Genoveva Martí, eds., *Empty Representations: Reference and Non-Existence*, 25–66. Oxford University Press.

Dorst, Kevin, 2019. 'Abominable KK Failures'. *Mind*, 128(512):1227–1259.

Dorst, Kevin and Mandelkern, Matthew, 2021. 'Good Guesses'. *Philosophy and Phenomenological Research*.

Edgington, Dorothy, 1995. 'On Conditionals'. *Mind*, 104(414):235–329.

Goodman, Jeremy and Holguín, Ben, 2022. 'Thinking and Being Sure'. *Philosophy and Phenomenological Research*, 1–27.

Goodman, Jeremy and Lederman, Harvey, 2021. 'Perspectivism'. *Noûs*, 55(3):623–648.

Hawthorne, John, Rothschild, Daniel, and Spectre, Levi, 2016. 'Belief is Weak'. *Philosophical Studies*, 173(5):1393–1404.

[16] I would like to thank Kenneth Black, Kevin Dorst, Caspar Hare, Justin Khoo, Jack Spencer, Bob Stalnaker, Eliot Watkins, Roger White and two anonymous referees for their helpful feedback and advice.

Hoek, Daniel, forthcoming. 'Minimal Rationality and the Web of Questions'. In Dirk Kindermann, Peter van Elswyk, and Andy Egan, eds., *Unstructured Content*. Oxford University Press.

Holguín, Ben, 2022. 'Thinking, Guessing, and Believing'. *Philosophers' Imprint*, 1–34.

Kratzer, Angelika, 1986. 'Conditionals'. *Chicago Linguistics Society*, 22(2):1–15.

Lewis, David, 1976. 'Probabilities of Conditionals and Conditional Probabilities'. *Philosophical Review*, 85(3):297–315.

Mandelkern, Matthew and Dorst, Kevin, 2022. 'Assertion is Weak'. *Philosophers' Imprint*.

McGee, Vann, 1985. 'A Counterexample to Modus Ponens'. *Journal of Philosophy*, 82(9):462–471.

Moss, Sarah, 2019. 'Full Belief and Loose Speech'. *Philosophical Review*, 128(3):255–291.

Nagel, Jennifer, 2021. *The Psychological Dimension of the Lottery Paradox*, 48–73. Cambridge University Press.

Rothschild, Daniel, 2020. 'What It Takes to Believe'. *Philosophical Studies*, 177(5):1345–1362.

Stalnaker, Robert, 1968. 'A Theory of Conditionals'. In Nicholas Rescher, ed., *Studies in Logical Theory (American Philosophical Quarterly Monographs 2)*, 98–112. Oxford: Blackwell.

———, 1980. 'A Defense of Conditional Excluded Middle'. In William Harper, Robert C. Stalnaker, and Glenn Pearce, eds., *Ifs*, 87–104. Reidel.

Williamson, Timothy, 2020. *Suppose and Tell: The Semantics and Heuristics of Conditionals*. Oxford, England: Oxford University Press.

———, forthcoming. 'Knowledge, credence, and the strength of belief'. In A. Flowerree and B. Reed, eds., *Towards an Expansive Epistemology: Norms, Action, and the Social World*. Routledge, London.