Counterfactual Skepticism is (Just) Skepticism

David Boylan*

Counterfactual skepticism says that most ordinary counterfactuals are false. Its devotees are few, but they have a simple, powerful argument in their arsenal: most counterfactuals are chancy; and chancy counterfactuals, they say, cannot be true.

The camp of the non-skeptics — much better populated — spends its time on defensive strategies.[1] Stories are given about how counterfactuals *could possibly* be true, notwithstanding the arguments of the counterfactual skeptic. The defenders presume they only need to hold out: if they can deflect the force of the skeptical argument, counterfactual skepticism will melt away; so heavy are the costs of counterfactual skepticism in their eyes.

Should the non-skeptics be so confident? Hajek, the arch counterfactual skeptic, certainly thinks not. He claims counterfactual skepticism has not just a good argument but an attractive world-view. It is the best way, he suggests, of reconciling what we know about counterfactuals with discoveries about fundamental physics. What's more, he claims, this skepticism is benign. The ordinary counterfactuals we assert and reason about can be replaced with more hedged, probabilistic counterfactuals. Perhaps we say false things all the time; but we still hew close enough to the truth.

I have two aims in this paper. My first and primary aim is add to

*Thanks to audiences at Texas Tech University, the 2022 Joint Session, the Berlin Language Workshop and the 2022 Dubrovnik Conference on Philosophy of Language, and to Melissa Fusco, Matt Mandelkern, Ginger Schultheis, Joel Velasco, and Snow Zhang.

[1]Note that Kocurek (2022) is a recent and admirable exception here. Kocurek shows that Hajek's core principle, *Chance Undermines Would*, is in tension with standard principles in the logic of counterfactuals. Kocurek sketches, but ultimately rejects, a skeptical semantics for counterfactuals that validates qualified versions of those standard principles. My arguments, by contrast, will target even the latter skeptical semantics.

the case that counterfactual skepticism is not benign. The world-view is much less attractive than it seems, I will argue, because it leads to significant skepticism about the future: if counterfactual skepticism is true, then we can have only very limited knowledge about the future. I give three arguments to this conclusion, taking as premises principles connecting knowledge of indicative conditionals to knowledge of disjunctions; and principles linking knowledge of indicative conditionals to knowledge of counterfactual conditionals. Such principles are no less obvious than those favoured by the counterfactual skeptic. And endorsing future skepticism is a much more serious cost than abandoning ordinary counterfactuals.

My second aim is to examine the consequences for non-skeptical theories. My arguments indicate that not only are many ordinary counterfactuals true, but that we also *know* them. How can this be reconciled with the very real tension between counterfactuals and chance? I consider two popular solutions, one purely epistemic and one involving contextualism about counterfactuals. Together they suggest something like contextualism about *knowledge* is required to fully resolve the skeptical threat.

1 Motivating Counterfactual Skepticism

I said counterfactual skepticism is the view that most counterfactual conditionals are false.[2] But what are ordinary counterfactuals? Here is an example:

(1) If I were to drop this plate, it would break.

Ordinary counterfactuals state dependencies that we rely on in all sorts of ordinary reasoning. The counterfactual skeptic claims most such claims are

[2]Why *most* and not *all*? Principally because certain counterfactuals are not undermined by Hajek's arguments. Hajek agrees that a counterfactual will be true when its antecedent entails its consequent. He also thinks probabilistic counterfactuals can be true. For instance, the probabilified variant of (1) is true:

(i) If I were to drop this plate, it would almost certainly break.

false.

Why be a counterfactual skeptic? Perhaps the most powerful argument, due like many others to Hajek (2022), focuses on chance.[3] Chance seems to undermine counterfactuals. Consider:

(2) # If I had entered the lottery, I would have lost; but if I had entered, there would have been a (very small) chance I would win.

How can I assert that I would have lost, if there is a chance I would have won? Even such a small chance undermines the first conjunct. The particular lesson Hajek draws is that (2) is false: the truth of a counterfactual $A \mathrel{\Box\!\!\rightarrow} C$ is incompatible with a non-zero chance of $\neg C$, had A been the case. That is, Hajek subscribes to the following principle:

Chance Undermines Would. $A \mathrel{\Box\!\!\rightarrow} [ch(\neg B) \neq 0] \models \neg(A \mathrel{\Box\!\!\rightarrow} B)$

The trouble is our world is much chancier than expected. (1) initially strikes us as impeccable. After all, if I drop a plate, the chances are extremely high that it will hit the ground and shatter. But our best physics suggests the chance is *merely* high — not certain. Various strange outcomes have a non-zero chance: if I drop the plate, rather than striking the floor, there is a non-zero chance it will quantum tunnel through it. Such events of course are astronomically unlikely. But they have some non-zero chance.

If *Chance Undermines Would* holds, then our initial judgement about (1) is mistaken. It is in fact false: for if I had dropped the plate, there is a chance that it would not have broken. Of course, there is nothing special about this particular example: all sorts of ordinary counterfactuals will be undermined by the same kind of argument, because all sorts of counterfactuals turn out to be chancy. So, the argument goes, most counterfactuals are false.

The path from *Chance Undermines Would* to counterfactual skepticism is extremely narrow; and so many non-skeptics will want to push back on

[3] Loewenstein (2021a) gives a different argument based on Sobel sequences. I believe her position is also targeted by the kinds of arguments I give below.

that principle. I will indeed eventually propose an alternative explanation of examples like (2). But for much of this paper, I will spot *Chance Undermines Would* to the skeptic so that we can better see where counterfactual skepticism leads. Still, the stronger the prima facie case for counterfactual skepticism, the clearer its interest.[4] So, before I truly spot the skeptic their principle, I will note two further points on their behalf.

The first is methodological. In response to (2), one might observe that not all clashes are the result of inconsistency: some, like Moore's paradox, are the result of some mere pragmatic inconsistency. However, the burden is usually on those who favour a pragmatic solution to give independent reason for pursuing that route.

To see this, consider an example. Knowledge is taken to be factive: if you know something, it is true. One argument for this principle is simply that denying factivity sounds incoherent:

(3) #I know it's raining but it isn't raining.

Now this observation does not strictly speaking rule out any pragmatic explanation. But reaching for a pragmatic explanation here would clearly be wrong headed. Why? Because we have no independent reason to reject factivity; nor do we have independent evidence for a pragmatic explanation.[5] A pragmatic explanation would simply be unmotivated.

There appears then to be a methodological room of thumb here: absent some independent evidence *against* a semantic explanation or *for* a pragmatic explanation, such clashes are good evidence for an inconsistency. By the end of the paper, having argued from counterfactual skepticism to epistemological skepticism, I will take myself to have given such independent evidence against a semantic explanation; but at this stage of the dialectic, I am happy to concede to the skeptic there is none.[6]

---

[4] Thanks to an anonymous referee for encouraging me to say more here.

[5] Notice that this is exactly where Moore's paradox is disanalogous: clearly there is great independent pressure to reject the claim that all truths are known; and third person attributions of the form ⌜p but she doesn't know p⌝ are clearly consistent.

[6] Some suggest counterfactual skepticism is itself independent evidence against a semantic explanation: they think it is an advantage of a semantic theory of counterfactuals

The second point on behalf of the skeptic, due to Hajek (2022), is that rejecting Chance Undermines Would generates puzzles of its own. If you reject Chance Undermines Would, then you allow for counterfactuals to be chancy, yet true. A natural thought, though, is that there should be limits to how chancy a true counterfactual can be: how could it be true, we might think, that a coin would land heads if tossed, if the coin is heavily biased towards tails?

But it is extremely hard to vindicate this thought. It is in tension with how risk agglomerates: the conjunction of many high chance events will often have a very low chance. Call this the problem of *constraining counterfactual chanciness*: if counterfactuals can be true, despite a low counterfactual chance to the contrary, then there is considerable pressure to think that they can be true, despite *extremely high* counterfactual chances to the contrary.

To illustrate, take Hajek's case where we have a coin weighted 0.9999 towards heads. If small chances are compatible with truth, it might be tempting to say the following is true:

(4) If I were to flip the coin 100 million times, it would land heads on the first toss.

After all, when we reject Chance Undermines Would, we are committed to saying chance is not inconsistent with counterfactuals. The same reasoning motivates the thought that the same is true for *any* toss: for any n it is true that:

(5) If I were to flip the coin 100 million times, it would land heads on the nth toss.

But it is extremely plausible that the consequents of counterfactuals agglomerate: if it is true that $A \mathrel{\Box\!\!\rightarrow} B$ and that $A \mathrel{\Box\!\!\rightarrow} C$, then it follows that

if it predicts that we mostly use them to speak truly; see Lewis (2016) for an example of this attitude. This of course begs the question against the counterfactual skeptic; and raises tricky dialectical questions about whether it is acceptable to beg this question against the skeptic. My arguments aim to beg no such questions.

$A \mathrel{\Box\!\!\to} (B \wedge C)$. Repeated applications of this rule will eventually give us:

(6) If I were to flip the coin 100 million times, it would land heads every time.

This might seem like an extremely unpalatable result: even with such a biased coin, it is *extremely* unlikely that it would land heads all of the 100 million times. We would be much more inclined to say that the contrary counterfactual is true:

(7) If I were to flip the coin 100 million times, it would *not* land heads every time.

After all, it is extremely likely that the coin would not land heads every time. Since their shared antecedent is possible, (6) and (7) cannot both be true.[7] But if (6) is false, then at least one instance of (5) must be false. Which one could that be? No particular instance of (5) seems apt for rejection; and denying some instances but not others seems unmotivated.

It is a tempting thought that a small amount of counterfactual chance is consistent with the truth of a counterfactual. What Hajek shows is that this is hard to contain: if you say that a a counterfactual can be true despite a small amount of the contrary, it is hard to resist the conclusion that it can still be true despite an *extremely high* chance to the contrary. A neat solution to this puzzle is to not admit any true, but chancy counterfactuals in the first place: a counterfactual cannot be true, if there is *any* chance to the contrary; *Chance Undermines Would* is valid.

2 The Qualitative Thesis

I give three arguments from counterfactual skepticism to future skepticism. All share a premise linking knowledge of indicative conditionals to knowl-

[7] Note that the qualification about the antecedent is necessary here: when A is not possible, then on the standard semantics for counterfactuals, $A \mathrel{\Box\!\!\to} C$ and $A \mathrel{\Box\!\!\to} \neg C$ will both be true.

edge of material conditionals. It has appeared under a few different names, but following Boylan and Schultheis (2022) I will call this premise the *Qualitative Thesis*.[8,9]

The Qualitative Thesis says that, when you leave open A and you know the material conditional $A \supset C$, you know the indicative conditional $A > C$:

Qualitative Thesis. If $\neg K \neg A$ and $K(A \supset C)$, then $K(A > C)$

To assess the Qualitative Thesis, recall that $A \supset C$ is equivalent to $\neg A \vee C$. Thus, according to the Qualitative Thesis, if for all I know a coin will be flipped and I know (either it won't be flipped or it will land heads), then I know that (if the coin is flipped, it will land heads).

The Qualitative Thesis, and its analogues for other attitudes like belief and certainty, are fundamental to our use of indicative conditionals. Ordinary reasoning very often treats the indicative conditional as interchangeable with the material. To see this, consider *the Direct Argument*:

The Direct Argument. $\Diamond \neg A, A \vee B \rightsquigarrow \neg A > B$

As many have observed, arguments of this form seem excellent:

(8) a. That card is either a club or a heart
b. (And it may not be a club.)
c. So if the card isn't a club, it's a heart.

In some sense, the Direct Argument is a good argument: generally we would indeed feel entitled to infer the conclusion from the premises.[10]

---

[8] See also Holguin (2021), who calls it Material Indication.

[9] What Boylan and Schultheis (2022) call the Qualitative Thesis is slightly stronger: it says that, when the antecedent is left open one knows $A > C$ *iff* one knows $A \supset C$. But the stronger thesis follows anyway from my other assumptions of Modus Ponens for $>$, classical logic and *Closure*.

[10] Notice it is indeed crucial that the antecedent of the conditional is left open. For consider the following example:

(i) a. Either it won't rain very hard or it won't rain at all.
b. $\not\rightsquigarrow$ So if it rains very hard, then it won't rain.

I say "in some sense" because there are different views on what that sense might be. A minority of theorists think the indicative conditional *is* the material conditional. For them, the status of the Direct Argument is no mystery: it is simply valid. That theorist will surely sign on for the Qualitative Thesis, assuming that knowledge is closed under known entailment:[11]

Closure. $K(A \supset C), KA \models KC$

The vast majority of theorists do not think the indicative is the material conditional. What then is their explanation of the Direct Argument? It must be, as Stalnaker (1975) originally observed, that the Direct Argument is a good argument in some *weaker* sense than classical validity. Stalnaker proposed that it is a reasonable inference: if one leaves open the left conjunct and one *accepts* the disjunction, then one should accept the corresponding indicative. Put slightly differently, the Qualitative Thesis, together with its analogues in terms of belief and certainty, provides an excellent explanation of why the Direct Argument strikes us as a good argument: it is *knowledge* preserving.[12]

Here the inference is not acceptable in any sense. But the natural explanation is that, if we accept (i), the antecedent of (ii) is not left open: (i) entails it won't rain very hard. For this reason, the Qualitative Thesis would not license this inference: that principle requires that the antecedent be left open. Nonetheless, very often the second premise is left tacit, as the assertion of the disjunction tends to carry the implicature that neither disjunct is known. (This implicature is crucial though for Stalnaker (1975)'s original explanation of the phenomenon.) Thanks to anonymous reviewer here for discussion.

[11] Both Closure and the Qualitative Thesis should be read as principles about being in a position to know, rather than knowing, if one is worried about cases where the requisite beliefs are missing. I continue to use the above formulations for simplicity.

[12] Given some other natural assumptions, the Qualitative Thesis also falls out of Stalnaker's Thesis:

> *Stalnaker's Thesis.* If Pr is a reasonable probability function, then $Pr(A > C) = Pr(C|A)$, when $Pr(A) > 0$.

The first assumption is that, if your prior Pr is reasonable, then so too is the result of updating it on the strongest proposition you know; call that updated probability function Pr_K. The second is that the propositions you know are exactly those that have probability 1, upon updating on what you know:

> *Knowledge is Pr_k 1.* KA iff $Pr_K(A) = 1$

Given the probability calculus, these three assumptions together suffice for the Qualitative Thesis.

The Qualitative Thesis is a fundamental constraint on indicatives. As we will now see, it pairs badly with counterfactual skepticism.

3 First Argument: Indicative Skepticism

My first argument focuses on indicatives. Hajek's style of argument equally well motivates an analogue of *Chance Undermines Would* for indicatives; but given the Qualitative Thesis, this leads to widespread *disjunctive skepticism*, that is, skepticism about the ordinary disjunctive claims corresponding to the counterfactuals Hajek targets. And disjunctive skepticism is only a short step away from skepticism about the future.

First, let's establish that chance also undermines indicative conditionals. Consider:

(9) # If the plate is dropped, it will smash; but if the plate is dropped there is a chance it will not smash.

This seems no better than (1), the example that motivated *Chance Undermines Would*. I take it then that the defender of counterfactual skepticism is committed to a similar principle about indicatives:

Chance Undermines Indicatives. $A > [ch(\neg B) \neq 0] \models \neg(A > B)$

Chance Undermines Indicatives leads to indicative skepticism, the claim that most ordinary indicatives are false, by the same style of argument as from §1.

Given the Qualitative Thesis, indicative skepticism leads to disjunctive skepticism. Suppose I am holding a plate, and wondering whether to let it drop. We said in §1 that if I do drop it, there is a non-zero chance it will take some weird trajectory and so not smash. *Chance Undermines Indicatives* leads us to the conclusion that the indicative is false:

Hajek of course will note that Stalnaker's Thesis has been subjected to many *triviality results*. But various forms of the Thesis have been shown to be tenable: Bacon (2015) proves that a contextualist version Stalnaker's Thesis, and hence the Qualitative Thesis, is tenable; Goldstein and Santorio (2021) demonstrate tenability for an invariantist approach to the indicative. My arguments will go through on either approach.

(10) $\neg$(the plate is dropped > it smashes)

But knowledge is factive; so we cannot then know that conditional:

(11) $\neg K$(dropped > smashed)

Now since the antecedent is clearly left open, by the Qualitative Thesis, it's false that we know the corresponding material conditional:

(12) $\neg K$(the plate is dropped $\supset$ it smashes)

And material conditionals are equivalent to disjunctions. So, given Indicative Skepticism, I do not know the equivalent disjunction:

(13) $\neg K$(the plate isn't dropped $\vee$ it will smash)

If most ordinary indicatives are false, then by the Qualitative Thesis, we cannot know the corresponding disjunctions.

How bad is disjunctive skepticism like this? Quite bad: it leads to widespread skepticism about the future. Consider some mundane cases of future knowledge:

> **Plate Freefall.** Suppose I can see a plate has just slipped out of your hand, but I don't know whether you dropped it with your left or your right hand.
>
> I know the plate is going to smash.
>
> **Going Home.** Suppose I alternate my route home between two equally long routes, A and B, both taking 30 min. Either way, I'm going to leave the office at 5.
>
> I know I will be home at 5.30.

Disjunctive skepticism is incompatible with future knowledge in either case, given Closure for knowledge.[13] Take **Freefall**. If I know the plate is going

[13] Notice all that is really needed here is the plausible principle that knowledge is preserved under *disjunction introduction*: if I know A then I know $A \vee B$. even those like Closure-deniers like Nozick (1981) tend to accept this weak instance of Closure.

to smash, then by Closure I know the disjunction:

(14) $\neg$(you dropped it with your left hand) $\vee$ the plate will smash.

Likewise in **Going Home** I know:

(15) $\neg$(I go home by route A) $\vee$ I will be home at 5.30.

Disjunctive skepticism says I can't know either of these disjunctions; by Closure, I can't know any disjuncts entailing them either. So disjunctive skepticism is incompatible with saying we have knowledge in these mundane cases. But of course, given the Qualitative Thesis, *Chance Undermines Indicatives* leads us to disjunctive skepticism.

There is a general recipe here for deriving future skepticism from *Chance Undermines Indicatives*. Take some ordinary future event E that is *slightly* chancy and yet known. Divide it into two subcases, C and C', neither of which you know will obtain and both of which allow a slight chance for E to be false. In such a case, if you know E, Closure says you should also know $\neg C \vee E$. Since $\neg C$ is compatible with your knowledge, the Qualitative Thesis tells us that you will also know $C > E$. But this is incompatible with *Chance Undermines Indicatives*: $C > E$ should be false, if there's even the slightest chance that E might happen to be false, while C is true.

This recipe yields a lot of future skepticism; and so the Qualitative Thesis give us a powerful reason to reject *Chance Undermines Indicatives*. Now of course the counterfactual skeptic is not necessarily wedded to *Chance Undermines Indicatives*. But *Chance Undermines Indicatives* is motivated by data that *precisely* parallels the data for *Chance Undermines Would*. What reason could there be to reject *Chance Undermines Indicatives* while embracing *Chance Undermines Would*?

4 Second Argument: Future Directed Counterfactuals

Given the first argument, we can mount two more direct arguments from counterfactual skepticism to future skepticism. The key is to exploit certain connections between our knowledge of counterfactuals and of indicatives. From those principles and the argument of the previous section, we can argue from counterfactual skepticism directly to future skepticism.

The first of these arguments focuses on future directed counterfactuals and indicatives. Say that a counterfactual is *future* directed (at a time of utterance) if its antecedent is about some time in the future (of the time of utterance). Let's start by observing that the following principle seems like a good one:

> *Future directed.* If at t you know a future directed indicative, then at t you know the corresponding future directed counterfactual.

Why think this? Because almost always, it sounds off to say things like:[14]

(16) # I know that the plate will smash, if it is dropped; but I don't know whether it *would* smash, if it *were* dropped.

(17) I know that I will get home in an hour, if I leave now; but I don't know whether I *would* get home in an hour, if I *were* to leave now.

A good *explanation* of this is that there seems to be very little room between forward-directed counterfactuals and indicatives. Many are tempted by the view that future indicatives and counterfactuals are in fact equivalent.[15] That being said, strictly speaking I need only the weaker principle above.

Given *Future Directed*, we can connect the arguments of §3 directly to counterfactuals. *Chance Undermines Would* applies just as well to future directed counterfactuals as any other. It sounds no better to say:

(18) If the coin were to be flipped it would land heads; but if it were flipped there's a chance it would land tails.

[14] In §6.1 I discuss potential counterexamples to this principle.

[15] Dudman (1994) argues that they are equivalent and Jackson (1990) gives a Sly Pete example in favour of this position. See also Lycan (2001) and DeRose (2010) for further sympathetic discussion.

So the counterfactual skeptic must say most ordinary future directed counterfactuals are false. But given *Future Directed*, we conclude that the analogous future directed indicatives are not known. Together *Future Directed* and *Chance Undermines Would* entail widespread epistemological skepticism about indicatives.

From here, we can simply plug in the argument from §3. Focus on future-directed counterfactuals where we do not know whether the antecedent holds: give counterfactual skepticism, the corresponding indicative is false and so we do not know it; and by the Qualitative Thesis, we do not know the corresponding disjunction. And so by Closure, we cannot know any disjuncts entailing such disjunctions. Running the argument in reverse, if we can have knowledge of slightly chancy facts, then we should know various forward directed indicatives and so should also know various forward directed counterfactuals. This runs contrary to counterfactual skepticism.

5 Third Argument: Robust Knowledge

Finally, I argue against skepticism about past directed counterfactuals.[16] This argument relies on a different principle linking counterfactuals and indicatives: very often, if I start off knowing an indicative, I would come to know the corresponding counterfactual, if I learned the antecedent were false. I argue this leads to skepticism in cases where our background evidence is *robust*, in a sense to be defined.

5.1 Prerequisites

First I'll motivate the principle and the definition we need to run the argument.

Return to our simple case of the plate. Initially I regard myself as knowing that if the plate is dropped at 2pm it will smash. But 2pm comes and

[16] Of course, for this strategy to get off the ground, the counterfactual skeptic would have to claim that past directed counterfactuals do not express the same thing as future directed counterfactuals at an earlier time. I will simply spot them this assumption.

goes and I see that the plate is not dropped. Now I regard myself as knowing that if it *had* been dropped at 2pm, it *would* have smashed. My knowledge of the indicative has turned counterfactual, as it were. This is commonplace: if I start off knowing an indicative, I often come to know the corresponding counterfactual, when I learn the antecedent is false.[17]

For further evidence, notice that various factors that would undermine my later knowledge of the counterfactual would do the same to my knowledge of the indicative. Suppose I learn the plate is indestructible. I will certainly give up the belief that it would have smashed, if it had been dropped; but I would also be inclined to think I did not know the indicative either.

Or suppose it turns out that my later belief in the counterfactual is Gettiered: it's true that the plate would have smashed; but unbeknownst to me, that plate is extremely robust and only would have smashed because only because God would have intervened to make it smash. In such a situation, it seems like my earlier belief in the indicative is Gettiered too.

Or suppose I learn the plate came from a factory, where half of the plates made are indestructible. I am clearly not even justified in believing the counterfactual in this situation; and I will think that had I known this earlier, I wouldn't have been justified in believing the indicative either.

Let's sum this up in a principle. Very often the following holds:[18]

> *Indicatives Turn Counterfactual.* If having total evidence E suffices for you to know $A > C$, then having total evidence $E \cap \neg A$ suffices for you to know $A \mathrel{\Box\!\!\rightarrow} C$.[19]

Theses like this are entailed by natural theories of the relationship between indicatives and counterfactuals. Those, like Adams (1970), who think that past counterfactuals are simply past tensed indicatives will endorse the principle. Others, like Edgington (1995, 2008), say that past counterfactu-

[17] Actually the principle will be neutral on whether we *come* to know the counterfactual. It simply says we *do* know it, after we learn the antecedent is false. For my purposes it doesn't matter whether we did not know it before we learned the antecedent is false.

[18] I discuss potential counterexamples to this principle in §6.1.

[19] I'm going to assume that one's evidence is knowledge, to state this principle in a natural way. Those who dissent should simply replace talk of evidence with talk of knowledge.

als talk about some *hypothetical* information state; and they will think that often that hypothetical state is often a past information state.

Crucially, this principle will enable me to extend the arguments to past directed counterfactuals. Very often we learn that the antecedent is false because we observe the event in question never happens: 2pm comes and the plate is not dropped. So I will focus on cases where we learn the antecedent is false after the relevant time has passed. This will allow us to see that counterfactual skepticism has skeptical consequences, even when we restrict our attention to past directed counterfactuals.

The other notion I need to define is that of one's background evidence being *robust* with respect to a proposition A.

For concreteness let's start with the following minimal pair of cases, based on the original plate case:

> *Plate Dropped.* Again, you have a plate that I know is fragile. Suppose that at 1.30pm, I learn that the plate will be dropped at 2pm. This is all I learn between 1 and 3pm; and my evidence at 1 pm is exactly the same as it was in *Plate Dropped*. The plate is indeed dropped at 2pm and smashes in a normal way.

> *Plate Not Dropped.* Again, you have a plate that I know is fragile; this time it is not dropped at all. Suppose that just after 2pm I learn that the plate was not dropped at 2pm; and this is all I learn between 1 and 3pm.

There's an important commonality here. Up to 1.30pm, my evidence is the same in both cases. Afterwards there is a divergence: in one case I learn a certain proposition, in the other I learn its negation. I'll say that my background evidence with respect to A is *robust* in w iff there are nearby worlds where the chances are the same, where I have exactly the same background evidence but I go on to learn $\neg A$ instead of A. (In §5.3, I make precise the notion of one's background evidence in terms of the contraction operation from the belief revision literature.) In *Plate Dropped*, my background evidence with respect to the proposition *the plate is dropped* — here just my earlier evidence that the plate is fragile— is robust, as witnessed by *Plate Not Dropped*: the latter is a nearby case, where prior to 1.30pm I know exactly the same things about the plate, but where the plate is not dropped.

5.2 Future Skepticism

We can now run the argument for future skepticism. First, I focus on the pair of cases above: I derive disjunctive skepticism in *Plate Not Dropped* in much the same way as before and use this to argue for future skepticism in *Plate Dropped*. This done, I will give the more general recipe for future skepticism in cases of robust background evidence.

By now, it's routine to establish disjunctive skepticism in *Plate Not Dropped*. Per counterfactual skepticism, the relevant counterfactual is false:

(19) $\neg$(plate dropped at 2pm $\Box\!\!\rightarrow$ plate smashed)

So I do not know it. But then it follows, by *Indicatives Turn Counterfactual* that I did not know the corresponding indicative earlier at 1pm:

(20) $\neg K$(plate dropped at 2pm > plate will smash)

But then, given the Qualitative Thesis, I did not at 1pm know the corresponding disjunction:

(21) $\neg K(\neg$(plate dropped at 2pm) $\vee$ plate will smash)

If the counterfactual is false, I cannot know it; then by *Indicatives Turn Counterfactual* I cannot earlier know the indicative in the weaker evidential state I was in earlier; and by the Qualitative Thesis I cannot know the corresponding disjunction in the earlier state.

But disjunctive skepticism in *Plate Not Dropped* leads to future skepticism in *Plate Dropped*, precisely because my background evidence is robust in the latter case. I stipulated that the evidence is the same at the beginning of both cases: our evidence is just the fact that the plate is fragile. As a consequence, (21) holds in *Plate Dropped* iff it holds in *Plate Not Dropped*: I know the relevant disjunction in both cases or neither. We just saw that it is not known in *Plate Not Dropped*. So it is not known in *Plate Dropped* either.

This is incompatible with knowing the plate will smash in *Plate Dropped*. If I don't know the disjunction above, then at the earlier time there are three kinds of worlds compatible with what I know: worlds where it is not dropped; worlds where it smashes; and worlds where it is dropped but does not smash. At 1.30pm in *Plate Dropped*, all I learn is that the plate will be dropped. So I now can eliminate the first kind of world; but there are still two kinds of world left: the ones where it is dropped and smashes; and the ones where it is dropped but does not smash. So I do not know that the plate will smash, all because I did not know the disjunction to begin with.

So we see that disjunctive skepticism has ramifications for other nearby cases. This particular pair of cases is illustrative, but again let's extract the more general recipe.

Suppose we have a case where my background evidence is robust with respect to A and where I learn that some slightly chancy future claim B is true on the basis of learning A. I claim that, given our assumptions, counterfactual skepticism leads to a contradiction.

Counterfactual skepticism gives us that $\neg A \vee B$ cannot be part of my background evidence E. For by robustness, we know there is a second case, where my background evidence is also E but I learn $\neg A$; since counterfactual skepticism says it is false there, I can't know $A \mathrel{\Box\!\!\rightarrow} B$ on the basis of $E \cap \neg A$; by *Indicatives Turn Counterfactual* I can't know the indicative $A > B$ on the basis of E; and, by the Qualitative Thesis, I can't know the disjunction $\neg A \vee B$ on the basis of E either. Since E is my background evidence in *both* cases, we get that in the first case I don't know $\neg A \vee B$ before learning A.

This disjunctive skepticism leads to future skepticism: if $\neg A \vee B$ is not part of my background evidence, then learning A will not enable me to know B: there will be worlds consistent with my evidence where $A \wedge \neg B$ that remain consistent with my evidence, even after I learn A. So, contrary to our supposition, I do not learn the future claim B after all.

This recipe would lead to a lot of future skepticism. It's just a fact of life that, in very many cases, whether we have future knowledge depends partly on what is entailed by our robust background evidence. It might even be the *most typical* way we get future knowledge. I might decide I am going to take my normal route home and so come to know that I will be home in 30 minutes. Here my background knowledge is robust: it consists of simple background facts about the layout of my neighbourhood and how quickly I drive, all things that are consistent with taking an alternate route home. To take another example, if I learn that Alice is going to jump, I know she will come down. Again my background knowledge is robust: it consists of knowledge of various dependencies and those dependencies would still have held, if I had learned Alice didn't jump.

5.3 Formalising the Argument

Before moving on, I will give a more formal version of the recipe for future skepticism.

To formalise my notion of background evidence, I use the notion of *contraction* from the AGM literature.[20] In AGM, contraction is an operation on a belief state where a proposition is removed from your beliefs, while making the minimum overall changes necessary to one's beliefs.

In the argument below, I will use two accessibility relations: $R(w)$, the set of worlds consistent with what one knows at w; $R^{-A}(w)$, the set of worlds consistent with one's knowledge at w contracted with A. In AGM these accessibility relations are assumed to obey the following constraints:

Success. If $A \neq \top$ then $R^{-A}(w) \nsubseteq A$

Recovery. $(R^{-A})(w) \cap A = R(w)$

Success says that your background evidence for A does not entail A, whenever A is a contingent proposition. Recovery says one's background evidence for A intersected with A yields one's entire evidence.

[20]See Alchourrón et al. (1985) and Grove (1988) here.

Contraction gives us a way to make sense of the notion of one's background evidence with respect to a particular proposition: it is simply their evidence, contracted with A. (Often (but not necessarily always) the contraction of one's evidence will simply be what they knew before they learned A.) On this understanding, one's background evidence with respect to A entails B iff $R^{-A}(w) \subseteq B$.

Given the contraction analysis of background evidence, we can define robustness as follows:

> $R^{-A}(w)$ is robust iff there's some nearby world w' where $R(w') \subseteq \neg A$ and $R^{-\neg A}(w') = R^{-A}(w)$.

I assume that a necessary condition on being nearby in this sense is that the counterfactual chances remain the same.

> *Nearby.* If w' is a nearby world to w, then $w \in (A \mathrel{\Box\!\!\rightarrow} Ch(C) = x)$ iff $w' \in (A \mathrel{\Box\!\!\rightarrow} Ch(C) = x)$.

Now we can spell out more precisely the argument in the text above. When our background evidence with respect to A is robust, we can derive a contradiction from the claim that $R(w) \subseteq C$, assuming that $A \mathrel{\Box\!\!\rightarrow} Ch(\neg C) > 0$:

1. $R(w) \subseteq C$ (Assumption)
2. $w \in (A \mathrel{\Box\!\!\rightarrow} Ch(C) > 0)$ (Assumption)
3. $\exists\, w'$ such that $R(w') \subseteq \neg A$ and $R^{-\neg A}(w') = R^{-A}(w)$. (Robustness)
4. $w' \in (A \mathrel{\Box\!\!\rightarrow} Ch(C) > 0)$ (*Nearby*)
5. $w' \in \neg(A \mathrel{\Box\!\!\rightarrow} C)$ (*Chance Undermines Would*)
6. $R(w') \nsubseteq (A \mathrel{\Box\!\!\rightarrow} C)$ (*Factivity of K*)
7. $R(w') = (R^{-\neg A})(w') \cap \neg A$ (*Recovery*)
8. $R^{-\neg A}(w') \nsubseteq A > C$ (*Indicatives Turn Counterfactual*)

9. $R^{-\neg A}(w') \nsubseteq \neg A$ (*Success*)

10. $R^{-\neg A}(w') \nsubseteq \neg A \vee C$ (Qualitative Thesis, 8,9)

11. $R^{-A}(w) \nsubseteq \neg A \vee C$ (from Robustness and 10)

12. $R^{-A}(w) \cap A \cap \neg C \neq \varnothing$ (11)

13. $R^{-A}(w) \cap A = R(w)$ (*Recovery*)

14. $R(w) \cap \neg C \neq \varnothing$ (12, 13)

15. $R(w) \nsubseteq C$ (14)

16. $\perp$ (1,15)

6 Responses

Future skepticism is too high a cost to bear, so the counterfactual skeptic must reject some step in my arguments. There are a number of strategies here: they might reject the principles linking indicatives and counterfactuals; or they might reject the Qualitative Thesis, either on dialectical grounds or in favour of a weaker chancy substitute. I argue all these strategies fail.

6.1 Reject the Linking Principles, First Pass

The counterfactual skeptic might respond by rejecting *Future Directed* and *Indicatives Turn Counterfactual*.[21] For both principles are subject to counterexample.

Here is a counterexample to *Future Directed*; the essential structure is Edgington (1995)'s.[22] Suppose I have good reason to think that either the butler or the gardener will tomorrow murder the Count. One of them harbours a bitter grudge against the Count and the other would never hurt a fly — I just don't know which is which. Here I know the indicative:

[21] Though note of course this would not help with my first argument.

[22] The literature contains other examples. Gibbard (1981) makes essentially the same point using a Sly Pete case. See also Khoo (2022) for different examples of believing indicatives without believing the corresponding subjunctives.

(22) If the butler doesn't try to kill the Count, then the gardener will.

But clearly I do not know the analogous counterfactual:

(23) If the butler were not to try to kill the Count, the gardener would.

Here I knew the indicative and subsequently learned the negation of its antecedent. Nonetheless, the counterfactual is clearly false.

A slight modification to the example yields a counterexample to *Indicatives Turn Counterfactual*; here I borrow from Bennett (2003) and Schulz (2017). Again I am trying to prevent the Count's murder and the setup is just as above. As before, I know the indicative:

(22) If the butler doesn't try to kill the Count, then the gardener will.

In fact the butler tries to murder the Count (but I foil him just in time). In retrospect, I say:

(24) If the butler hadn't tried to kill the Count, the gardener would have.

Here I knew the indicative and subsequently learned the negation of its antecedent. Nonetheless, the counterfactual is clearly false.

My argument does not require these principles to be exceptionless. If *Future Directed* and *Indicatives Turn Counterfactual* hold in a large range of cases, then counterfactual skepticism gives rise to a lot of future skepticism. For this reason, I believe the counterfactual skeptic cannot simply rest on these counterexamples.

What's more, plausibly some restricted versions of the principles are true. And it seems clear that those restricted principles should at least apply to the kind of cases we have been discussing so far. A first pass, suggested by Schulz (2017), is that the counterexamples to *Future Directed* and *Indicatives Turn Counterfactual* arise because there is a merely evidential connection between antecedent and consequent. In our examples, we do not imagine the butler and the gardener are cooperating: neither of their actions depend on the other. So we might simply restrict our principles to cases where this does not hold:

> *Restricted FD.* If at t you know a future directed indicative $A > C$ and that knowledge is not based on a merely evidential connection, then at t you know the corresponding future directed counterfactual $A \mathrel{\Box\!\!\rightarrow} C$.
>
> *Restricted ITS.* If having total evidence E suffices for you to know $A > C$ and that knowledge is not based on a merely evidential connection, then having total evidence $E \cap \neg A$ suffices for you to know $A \mathrel{\Box\!\!\rightarrow} C$.

I will not attempt to define the notion of merely evidential connection (though plausibly we could analyse it in terms of my notion of robustness). But I can point to cases where it is clearly absent. I take it to be incompatible with certain kinds of knowledge of counterfactual chance. If I know $\neg A \vee C$ and I also know that, on the supposition that A the chance of C is extremely high, then my knowledge of the disjunction is not merely based on an evidential connection. Further, I take this to transmit through the Qualitative Thesis: if I know $\neg A \vee C$ and I also know that, on the supposition that A the chance of C is extremely high, I know $A > C$ and not merely because of an evidential connection between antecedent and consequent.

With these restrictions in place, we can rerun the argument on the main cases from before. In the plate cases, I know that the chances of smashing are very high, supposing the plate were dropped. So if counterfactual skepticism holds, the only way I can know the indicative *dropped* > *smashes* is if my knowledge is based on a merely evidential connection, given *Restricted FD*. But then the only way I could know the disjunction ¬*dropped* ∨ *smashes* is if that knowledge is based on a merely evidential connection between the disjuncts. This is clearly wrong in the cases we have focussed on: if I know the disjunction, I know it on the basis of dependencies that hold in the world, the ones reflected in the chances. So it must be that I do not know the disjunction after all. From here the argument is the same as before.

6.2 Reject the Linking Principles, Second Pass: Partial Skepticism

An anonymous reviewer suggests a different route to rejecting the linking principles: we might adopt a skeptical view which treats indicatives very differently from counterfactuals, one where ordinary counterfactuals are false but ordinary indicatives are simply unknown, at least in high standards contexts. I call this view *partial skepticism*.

Here is the partial skeptic's picture in more detail. Suppose we accept counterfactual skepticism; what should we say about indicatives? I argued the counterfactual skeptic should also be an indicative skeptic by appeal to pairs like the following:

(9) #If the plate is dropped, it will smash; but if the plate is dropped there is a chance it will not smash.

But perhaps this is too quick. An alternative story is that this is merely a pragmatic, Moorean clash. A natural idea is that "knows" picks out different relations in different contexts. When we are ignoring weird possibilities (like brains in vats-worlds or quantum-tunnelling worlds), "knows" picks out a less demanding propositional attitude, $knows_{LO}$; and one can stand in the $knows_{LO}$ relation to a proposition A even if one cannot rule out weird and ignored possibilities where A is false. But once those possibilities become salient, "knows" picks out a more demanding attitude, $knows_{HI}$; to stand in the $knows_{HI}$ relation to a proposition, we must rule out even the weird possibilities where it is true.[23]

This kind of contextualist thought can be leveraged into an explanation of (9). We might say that ordinary indicatives can only be known in a low standards context, where weird possibilities such as quantum tunnelling are not salient; and once they become salient, we are pushed into a high standards context where ordinary indicatives are not known. What's more, even mentioning these possibilities tends to make them salient and so push

[23] I will in fact recommend something like this position for both indicatives and counterfactuals in §7.1.

us into a context where we do not count as knowing ordinary indicatives. Put together these pieces furnish an pragmatic explanation of (9): the right conjunct tends to push us into a context where the conjunction as a whole is not known.

The partial skeptic endorses this kind of explanation for indicatives; but maintains *Chance Undermines Would* and so endorses counterfactual skepticism. This is why their skepticism is partial: they only take most ordinary counterfactuals, not ordinary indicatives, to be false.

This view is important because it has the resources to answer my first two arguments. For the partial skeptic, my first argument fails at the first step: the partial skeptic rejects indicative skepticism and so the argument fails to get going. In my second argument, they will reject my use of the linking principle *Future Directed*. Partial skepticism entails *Future Directed* is false: in many low standards contexts, a future indicative is known; but, since ordinary counterfactuals are false, the corresponding counterfactual is not.

To argue for *Future Directed*, I noted that it sounds contradictory to deny it. We noted the badness of pairs like:

(16) #I know that the plate will smash, if it is dropped; but I don't know whether it *would* smash, if it *were* dropped.

Why would this be, if Future Directed is false? But, as an anonymous reviewer notes, the partial skeptic can give a pragmatic explanation where it merely appears to be valid.

It's plausible that often we speak falsely, but loosely. When I say

(25) It's 1 o'clock.

perhaps what I say is often *strictly speaking* false; but close enough to the truth to be assertable. The counterfactual skeptic can say the same about counterfactuals: most counterfactuals are false; but, when we are speaking loosely, ordinary counterfactuals are *close enough* to true in many contexts. Furthermore, it's plausible that there is a tight connection between the looseness of our assertions and what possibilities are (not) salient: when quantum tunnelling worlds are not salient, we are happy to speak loosely; but once they become salient, we become more strict.

The partial skeptic now has enough to explain the badness of (16). Suppose we are in a low standards context. Here I count as knowing the indicative. But, since we are speaking loosely, it is not appropriate to say that I do not know the counterfactual: it is approximately true and approximately true that I know it. Now suppose that we are in a high standards context. It is indeed appropriate to say that I do not know the counterfactual, since we are now speaking strictly. But it is not appropriate to say that I know the indicative: I do not count as knowing it by high standards.

Partial skepticism promises an interesting response to my arguments, if it is tenable. Should the counterfactual skeptic embrace this position? In fact, there are two important reasons to reject this view; these reasons I think should push everyone, counterfactual skeptic or not, towards rejecting partial skepticism.

The first reason is that the partial skeptic's combination of skepticism about counterfactuals but not indicatives is dialectically unstable: I know of no motivation for counterfactual skepticism which does not generalise to indicatives. Both of the motivations from §1 push us towards indicative skepticism. I noted a methodological rule of thumb: pragmatic explanations of clashes like (16) require independent motivation. If, in making the case for *Chance Undermines Would*, the counterfactual skeptic wants to lean on this point, then it should push them towards indicative skepticism too. I also sketched how those who reject *Chance Undermines Would* face the problem of constraining counterfactual chance. But it is clear that an analogous puzzle poses an equally strong case for *Chance Undermines Indicatives*: simply replace the counterfactuals everywhere with the corresponding indicatives.

The second reason is that a key commitment of partial skepticism is independently highly unintuitive. The partial skeptic accepts that indicatives may be unknown but possibly true, while counterfactuals are false: just suppose that we are in a high standards context, where quantum tunnelling possibilities are salient. But when it comes to future directed conditionals, specific examples show that this combination seems contradictory. Consider the following:[24]

(26) # I don't know whether if the coin is tossed it will land heads; but it's not the case that if it *were* tossed it *would* land heads.

(27) # I don't know whether if I drop the plate it will smash; but it's not the case that if it *were* dropped, it *would* smash.

Both of these claims sound inconsistent.

Now just as they did for *Future Directed*, the partial skeptic might try to give a pragmatic explanation here too, appealing to contextualism about knowledge and loose talk. But that explanation in fact does not in fact explain (26) and (27). Focus on (26) and suppose we are in a high standards context, where quantum tunnelling possibilities are salient. The left conjunct should be assertable, given the partial skeptic's view: their view says ordinary conditionals are not knowable in high standards contexts. But the right conjunct of (26) should also be assertable: after all, in our explanation of (16), we presupposed that raising such possibilities to salience forces us to speak strictly. So the partial skeptic's view in fact predicts that (26) and (27) should be perfectly fine things to say. They are not.

No further pragmatic story is obvious to me here. The markedness of

[24]For more evidence that this combination is contradictory, we can consider other attitudes. Consider the case of hoping: in order to hope that A, it seems one must neither know A nor $\neg A$. Now consider the following claim:

(i) # I hope that if the cup is dropped it will shatter; but it's false that it would shatter if it were dropped.

Again this sounds defective to my ear. Moreover, it sounds defective even when explicitly attending to quantum tunnelling possibilities:

(ii) # I hope that if the cup is dropped it will shatter; but of course because of quantum tunnelling it's false that it would shatter if dropped.

The natural explanation of this is that (i) and (ii) entail or commit a speaker to the truth of (26); but (26) is contradictory and so (i) and (ii) are defective. Parallel problems are posed by attitudes like wondering or fearing.

claims like (26) tell heavily against partial skepticism, I submit.

6.3 Reject the Qualitative Thesis on Dialectical Grounds?

One might worry the Qualitative Thesis is simply not dialectically effective in this context. Imagine the counterfactual skeptic arguing as follows. You've really just shown that counterfactual skeptics should also be indicative skeptics: if we accept *Chance Undermines Would*, then *Chance Undermines Indicatives* is not far behind. So, the skeptic might say, the Qualitative Thesis is question-begging here — for you've shown me, the counterfactual skeptic, that most indicatives are false too!

But this argument is misguided. The Qualitative Thesis does not by itself commit one to knowledge of indicatives. It is a conditional claim: one knows an indicative *if* one knows the corresponding material (assuming the antecedent is left open). This thesis is perfectly consistent with indicative skepticism. Of course it is *not* consistent with indicative skepticism plus the rejection of widespread disjunctive skepticism. But inconsistent triads are not automatically question-begging to those who accept two of the principles. The counterfactual skeptic would have to say that the *motivation* for one of the principles begs the question against their view.

But again, this is not the case. The motivation for the Qualitative Thesis is that the following is a good argument, in some sense of "good argument":

(28) a. Either the coin will not be flipped or it will land tails.
b. $\rightsquigarrow$ So if the coin is flipped, it will land tails.

In the same vein, it sounds just incoherent to assert either of the following:

(29) Either the coin will not be flipped or it will land tails; but if it is flipped it might land heads.

(30) Either the coin will not be flipped or it will land tails; but who knows if it will land tails if it is flipped.

I think we should understood this sense of "good argument" in terms of knowledge (or acceptance more broadly). Even so, we do not beg the question against counterfactual skepticism. These examples simply show that certain *combinations* are incoherent: if one knows $p \supset q$, one also knows $p > q$; one cannot know $p \supset q$ while being ignorant of $p > q$. Noting this kind of structural incoherence does not *by itself* commit us to saying we have any knowledge of indicatives.

Arguments for external world skepticism provide a helpful analogy. A simple skeptical argument uses the Closure principle from §2: I don't know whether I'm a brain in a vat; from Closure, if I don't know whether I'm a brain in a vat, I don't know whether I have hands; so I don't know whether I have hands. The use of Closure here is dialectically unproblematic.[25] But probably the best argument for Closure comes from noting that claims like the following are incoherent:[26]

(31) I know I have hands but I don't know whether I'm a brain in a vat.

Even a skeptic could use this as evidence for Closure, though they would deny that anyone knows they have hands. This is consistent for precisely the same reason as above: at this stage the skeptic is only ruling out certain *combinations* of knowledge and ignorance. Non-skeptics cannot reject the above argument by claiming that closure begs the question. The indicative skeptic cannot reject Qualitative Thesis in this way either.

6.4 Reject the Qualitative Thesis in Favour of a Substitute?

A different approach is to claim that our intuitions about the Qualitative Thesis really track a different principle. Counterfactual skeptics like to say that much of our ordinary counterfactual talk is loose. Really all we are in a position to assert are probabilistic counterfactuals like:

(32) If the plate were dropped, it would almost certainly smash.

[25] That is not to say Closure is completely unobjectionable. Rather it is to say the argument does not beg the question or suffer some other more subtle dialectical issue.

[26] See for instance DeRose (1995) and Hawthorne (2003).

As Hajek puts it, our ordinary counterfactual talk involves rounding errors. By analogy, one might think that the Qualitative Thesis also involves a rounding error. Perhaps the true principle is something like this:

High Chance QT. $\neg K \neg A, K(A \supset B) \models K(A > [Ch(B) \text{ is high}])$.

When we infer the simple indicative, we really add an extra rounding error to the conclusion of the High Chance QT.

I see no evidence that the Qualitative Thesis relies on a rounding error. Hajek's position is not unreasonable for simple counterfactuals: with the right pressure, the kind supplied by Hajek's cases, the person on the street will tend to retreat to things like (32). By contrast, a similar retreat to the High Chance QT is not in evidence, even when we raise to salience the chances of weird outcomes. The following is always defective to assert:

(33) Either the plate won't be dropped or it will smash; so if the plate drops, it's very likely *but not certain* that it will smash.

Assertions of (33) simply sound contradictory. Our intuitions about the disjunction and the indicative stand and fall together: if you make me retreat to a probabilistic indicative, you will thereby make me retreat from the disjunction too.

Putting it slightly differently, we would expect the following to sound coherent, if the Qualitative Thesis is not valid:

(34) Either the plate won't be dropped or it will smash; but if it is dropped it might not smash.

But it is not coherent, not even if it is salient that there is a chance it will not smash, conditional on being dropped.

Furthermore, we can control for our tendency to become hesitant about the disjunction by considering cases where we simply *suppose* it to be true. Consider the following discourse:

(35) a. In fact, if the plate is dropped, there is a small chance it will not smash because it might quantum tunnel.

b. But suppose as a matter of fact, it *is* the case that the plate won't be dropped or it will smash.
c. Then if the plate is dropped, it will smash.

This conclusion is extremely natural, much more natural than what the High Chance QT would yield. This is all the more impressive, given that the first line still highlights the chanciness of the plate smashing. Here of all places we should expect a retreat to the High Chance QT. We do not see it.

I conclude our intuitions do not track the High Chance QT; they track the Qualitative Thesis.[27] The Qualitative Thesis is *at least* as plausible as *Chance Undermines Indicatives*, but is not co-tenable given our knowledge of the future.

7 Non-Skeptical Responses

To show future skepticism is a cost *specific* to counterfactual skepticism, non-skeptical alternatives must do better. Here I consider two.[28] The first

[27]One further issue is that the High Chance QT is in serious tension with multi-premise closure. As Hawthorne and Lasonen-Aarnio (2009) and Williamson (2009) discuss, if we combine knowledge of the future with multi-premise closure, we tend to end up with knowledge of propositions that have a high chance of falsehood. Most future events are at least very slightly chancy; so non-skepticism about the future basically requires us to say that we can know a proposition, even if there is a small chance of it being false. Now suppose I know very many such propositions: by multi-premise closure, I must also know their conjunction. But the chance of their conjunction will be low, if they are conditionally independent.

Such conjunctions cause trouble for High Chance QT. Let C be such a conjunction which is known but has a low chance of truth and let $\top$ be an arbitrary tautology. By Closure, we know the disjunction $\neg\top \vee C$. By the High Chance QT we should then know that ($\top >$ high chance that C). But it is hard to maintain that this conditional is *true*, let alone known: the chance of C is very low.

This point rests on a controversial feature of multi-premise closure, namely knowledge despite high chance of falsehood. That being said, all views about this problem are forced to say something very unintuitive. To many, multi-premise closure seems worth the cost: it is a consequence of many of the most developed formal models of knowledge. For the defender of the High Chance QT it should be unwelcome news that those models are in tension with their principle.

[28]I don't mean these to be exhaustive; though I think it is particularly clear both how these views accommodate both future knowledge and the motivations for counterfactual

is an epistemic response, which rejects counterfactual skepticism wholesale. The other is a contextualist response, which grants that counterfactual skepticism is holds, but only in special skeptical contexts. An important lesson emerges from both discussions: responses to *epistemological* skepticism play a key role in heading off the skeptical threat.

7.1 The Epistemic Response

The epistemic response says that no counterfactuals are made false by chance alone. Most ordinary counterfactuals are both slightly chancy and true. The perceived tension between counterfactuals and chance is really a tension between *knowing* a counterfactual and chance. We do know, in the ordinary sense, various counterfactual claims. But when low probability events are raised to salience, we become more reluctant to attribute knowledge of those claims. Versions of this response have been given by Hawthorne (2005) and Stefánsson (2018).[29]

To motivate the epistemic response, start with an analogous puzzle about chance, knowledge and the future, loosely based on a puzzle from Hawthorne and Lasonen-Aarnio (2009). Take an example like the following:

---

skepticism.

[29] Moss (2013) and Kocurek (2022) also offer accounts in a similar vein, but with more significant differences of detail.

Kocurek endorses the parallel between future claims and counterfactuals, but explains the tension in terms of historical validity: claims like (36) are historical, but not classical, contradictions. It is not clear to me, however, that this suffices to explain their status. Given Kocurek's account of historical validity, it is not obvious that historical contradictions cannot be known; and if they *can* be known, it is not obvious why they should not be assertable. For the same reason, it is not clear to me whether this theory fully explains away the appeal of *Chance Undermines Would*.

While much of her discussion is sympathetic to how I sketch the view above, one worry about Moss's position is that it is unclear it allows for us to affirm that ordinary counterfactuals are known. Her central thesis is semantic humility, that our theories of the counterfactual shouldn't take a standard on whether ordinary counterfactuals are *actually* true. I worry that this stance of humility might also be inconsistent with a commitment to knowledge of the future, by a similar style of argument to the ones above. But I will not develop this worry in detail here.

(36) The plate will smash, but there is a small chance it will quantum tunnel.

This is no more assertable than the examples motivating *Chance Undermines Would*. But (36) cannot be classically inconsistent. If it was, that would mean a chance of tails entails that the coin *will* land tails — an absurd result that no theory of chance should deliver. It is more plausible that (36) is a Moorean contradiction. It is unassertable because it cannot be known: one cannot know A, we might think, if (one knows) there is a non-zero chance that $\neg A$.

That last claim is a principle linking knowledge and chance. Plausible as it might seem, it is a skeptical principle. We know now that our world is chancier than we thought. The chance/knowledge principle would say that we have thereby discovered we know a lot less about the future than we thought. This thought should be unwelcome, certainly more unwelcome than a rejection of the chance/knowledge principle. But something like the chance/knowledge principle seems *required* to explain why claims like (36) are not assertable.

That is the puzzle. But the puzzle has an attractive solution. Plausibly, what we take ourselves to know depends on what possibilities we are ignoring. And mentioning events like quantum tunnelling tends to *prevent* us for ignoring them. The contextualist view of knowledge, which we saw in the discussion of partial skepticism, is one way to spell this out; I frame the epistemic view in contextualist terms going forwards.[30,31]

The response then to our initial puzzle is that the chance/knowledge

[30] See Lewis (1996) for the canonical formulation.

[31] An alternative strategy here is the model of assertion from Moss (2012). On Moss's view, assertion must avoid *epistemic irresponsibility*, which Moss defines as follows:

> It is epistemically irresponsible to utter sentence S in context c if there is some proposition ϕ and possibility μ such that when the speaker utters S:
>
> 1. S expresses ϕ in C
> 2. ϕ is incompatible with μ
> 3. μ is a salient possibility
> 4. the speaker of S cannot rule out μ.

principle holds for high-standards knowledge, but not low-standards knowledge. I know$_{LO}$ that the plate will smash, despite the chance it might not: non-zero chance of falsehood does not immediately prevent knowledge. I do not know$_{HI}$ the plate will smash because that does require there to be a chance of 1 that it will smash. Finally, claims like (36) are not assertable because they tend to push us into high standards contexts; and in those contexts, we cannot be said to know (36).

The epistemic response simply applies this model to the problem of counterfactual skepticism. It rejects entirely the counterfactual skeptic's thesis: most ordinary counterfactuals are true. The relationship between chance and counterfactuals is precisely parallel to that between chance and future claims. Counterfactuals are not immediately undermined by chanciness: many are true even when they have a chance of falsehood. What's more, a certain grade of knowledge of counterfactuals is not undermined by chanciness. We can have ordinary, low-standards knowledge of counterfactuals that have a chance of falsehood. What we cannot have is high-standards knowledge of such a counterfactual; just as with future claims, a counterfactual cannot be known, if it has a non-zero chance of falsehood.

The epistemic response gives a simple explanation of where my argument goes wrong: counterfactual skepticism, the first premise, is false. Even still, it explains the motivation for *Chance Undermines Would*. It is predicted that we cannot assert claims like (9), repeated below:

(9) If the plate is dropped, it will smash; but if the plate is dropped there is a chance it will not smash.

The explanation is just as it was for (36): mentioning the chance of the counterfactual being false raises such possibilities to salience; this renders the

The Mossian explanation of our puzzle is similar, but goes via epistemic irresponsibility: normally it is fine to assert claims about the future, since weird low chance possibilities are not salient. But (36) is never acceptable because it *always* makes those claims salient.

This is of course a different solution to our puzzle. But just like contextualism it makes crucial use of the notion of salience. Thus I think defenders of the Mossian view should agree with the broader moral I draw in this section.

claim unassertable. But clearly this does not amount to validating *Chance Undermines Would*. This contextualist Moorean explanation allows most ordinary counterfactuals to be true, while still explaining the appeal of Hajek's principle. Thus we reject my arguments at the very first step, while explaining their appeal.

But what semantics would allow for chancy, yet true counterfactuals? Hawthorne (2005) argues persuasively that *similarity* analyses of counterfactuals will not do. I agree with Hawthorne that a variant on the Stalnakerian analysis a good candidate for the job. On Stalnaker's semantics for the counterfactual, $A \,\Box\!\!\rightarrow C$ is true just in case the selected A-world is a C-world. The obvious question is *how* such a world gets selected in the first place. I recommend a recently emerging picture from Schulz (2014) and Bacon (2015), where selection is understood, not in terms of similarity, but really in terms of *random* selection: each world *randomly selects* a closest A-world; and $A \,\Box\!\!\rightarrow C$ is true at w iff the randomly selected A-world at w is a C-world.

This picture makes the analogy between counterfactuals and the future extremely tight. How can we know facts about the future, even though it is sometimes chancy? Because the danger of being wrong is low: the worlds where, for instance, the plate quantum tunnels to China are very unlikely; and it would be very abnormal to be in such a world. How can we know counterfactuals, even though they are chancy? The answer is almost exactly the same: the chances are very low that a quantum-tunnelling world would be the selected world where the plate is dropped; and it would also be very abnormal for our world to select a world like that. For this epistemic Stalnakerian picture, the two issues very clearly become one and the same.[32]

[32]This analogy between counterfactuals and the future also helps see how the epistemic Stalnakerian deals the problem of constraining counterfactual chanciness. They simply accept that a counterfactual can be true despite an arbitrarily high counterfactual chance to the contrary; but their view can make sense of why this would be.

First of all, note that, when it comes to future claims, there is no problem about constraining chances. Return to Hajek's heavily biased coin. Suppose that in fact it *will* be tossed 100 million times. We know the following has extremely low chance:

I think this speaks very strongly in favour of the epistemic solution. It reduces the challenge of counterfactual skepticism to a more general skeptical challenge, one that we have the tools to solve. But this will be less impressive to those who find a Stalnakerian semantics is independently implausible. Many worry about the fact that the Stalnakerian validates the principle of Conditional Excluded Middle:

(i) The coin will land heads on every toss.

But that is perfectly consistent with its being true. For (i) to be true, all that needs to happen is that the actual future is indeed in the tiny minority of all heads worlds; and this may indeed be the case, even if it is extremely likely. So there is no problem of constraining chances of the future: a claim about the future may turn out to be true, despite having an arbitrarily low chance.

The epistemic Stalnakerian takes the same relationship to hold between chance and counterfactuals. What would happen if a coin was tossed 100 million times? That depends on what happens in the randomly selected world for that antecedent. In all candidates for the selected world, there is an extremely low chance it will land heads every time. But it still may happen that in that world the coin comes up heads every time. If indeed such a world is selected, then it is true that if the coin were tossed 100 million times, it would land heads every time. We should not *expect* this to happen: out of all the candidates for the selected world, only a tiny minority will be ones where the coin does this. But for the Stalnakerian this is nonetheless perfectly consistent: just as the actual future might be one out of some tiny minority of worlds with a very low chance, so too the selected world might be selected out of a similar set of low chance worlds.

This of course does not mean that this counterfactual is assertable: even if it were true that if the coin were tossed 100 million times, it would come up heads every time, the chance is so low as to make it never assertable. So here too their diagnosis of our discomfort is that the conclusion of Hajek's argument is unassertable, rather than untrue. Remember that we have to make this move in the case of the future. (i) could also turn out to be true after all; but before it has actually happened and been seen to happen, (i) will never be assertable.

Hajek's coin does still leave a residual puzzle about *knowledge*. I have said that we can know counterfactuals, despite a chance of falsehood. This leaves open that we might know, for each n, that the coin would land heads on the nth toss, if it had been tossed 100 million times. But then, assuming knowledge is closed under multi-premise deduction, we could know that the coin would land heads every time, if it had been tossed 100 million times.

But again a similar challenge arises for everyone in the case of the future: Hawthorne and Lasonen-Aarnio (2009) have shown how knowledge of the future generates a structurally similar puzzle. As we already saw, on pain of skepticism, we must admit that knowledge of the future is consistent with some chance to the contrary. But here too it is hard to maintain that knowledge is only consistent with *low* chances: given multipremise closure, knowledge of many propositions which individually have high chances requires knowledge of their conjunction, which will tend to have low chance. This puzzle about the future is a puzzle for everyone; and however we solve that puzzle will give the epistemic Stalnakerian a response to Hajek's coin.

Conditional Excluded Middle $A \mathrel{\Box\!\!\rightarrow} C \vee A \mathrel{\Box\!\!\rightarrow} \neg C$

A worry arises too about Humean supervenience: for the Stalnakerian, couldn't the non-conditional facts remain the same while a *different* world is the randomly selected world?[33]

I think that, on balance, the cost is worth the benefits: Counterfactual Excluded Middle has recently been given a number of powerful defences;[34] and, by appeal to techniques familiar from the literature on vagueness, the threat to Humean Supervenience can be reduced.[35] But this ultimately not my thesis. My aim is to show that there are a number of viable non-skeptical options that make clear how to avoid future skepticism. The epistemic response is one of them.

7.2 Contextualism

Contextualism about counterfactuals is a different reaction to Hajek's arguments, one pursued by Lewis (2016). I focus on Lewis (2016)'s *connective* contextualism, where it is the counterfactual connective itself that is context-sensitive.[36]

The contextualist says that the counterfactual skeptic is partly right: there is *a* reading of counterfactuals that is undermined by chance. Following Lewis (2016), it's helpful to assume a Lewisian variably strict account of the conditional: $A \mathrel{\Box\!\!\rightarrow} C$ is true at w iff the closest A-worlds — and there may be more than one — are C-worlds. The contextualist says that what p-worlds are closest is partly determined by *relevance in context*: in a world

[33]Hájek (2020) raises precisely this worry.

[34]See for instance Mandelkern (2018).

[35]For instance, Stalnaker (1981) gives a supervaluational defence of CEM, where a counterfactual is true in a context iff it is true relative to all Stalnakerian selection functions admissible in that context. Of course, many different random selections will be admissible; and so a large variety of selection functions will be admissible at every world. CEM is valid, since it is valid on all Stalnakerian selection functions. And there is no obvious threat to Humean Supervenience: at any world, lots of different random selections are compatible with our use of language; and so no *one* selection function picks out the counterfactual facts.

[36]I believe the same questions face the antecedent contextualism of Steele and Sandgren (2020). (I borrow the name from Loewenstein (2021b).)

where quantum tunnelling is not relevant, such worlds will not be closest; but they may be relevant once such possibilities become discussed.

This allows the contextualist to distinguish between two operators. When low probability events like quantum tunnelling are relevant, the counterfactual expresses the operator $\square\!\!\rightarrow_{c_{HI}}$; and when low probability events like quantum tunnelling are not relevant, the counterfactual expresses the operator $\square\!\!\rightarrow_{c_{LO}}$. The contextualist distinguishes between two versions of *Chance Undermines Would*:

Low Standards CUW. $A \square\!\!\rightarrow_{LO} (ch(\neg B) \neq 0) \models \neg(A \square\!\!\rightarrow_{LO} B)$

High Standards CUW. $A \square\!\!\rightarrow_{HI} (ch(\neg B) \neq 0) \models \neg(A \square\!\!\rightarrow_{HI} B)$

The contextualist accepts the second but denies the first. This accounts for the data motivating *Chance Undermines Would*, while avoid complete counterfactual skepticism. Merely mentioning events like quantum tunnelling will tend to push us into contexts where those possibilities are relevant; in those contexts the propositions expressed by counterfactual sentences are false, given *High Standards CUW*. But most of the time, this is not what we express by asserting counterfactual sentences. We assert propositions about $\square\!\!\rightarrow_{LO}$ and those are not threatened by low chance outcomes, since *Low Standards CUW* fails.

Simple contextualism does not immediately block my argument. Rather it faces two versions of my argument, one in terms of $\square\!\!\rightarrow_{LO}$ and another in terms of $\square\!\!\rightarrow_{HI}$. They reject the first argument at the first step; but what should they say about the second? After all, they say that the counterfactual skeptic is *right* about $\square\!\!\rightarrow_{HI}$-claims; they accept High Standards CUW.

Let's focus on my first argument. The contextualist will also distinguish between $>_{c_{HI}}$ and $>_{c_{LO}}$, the indicatives expressed in those high and low standards contexts respectively. The issue is that there does not seem to be *any* context in which the Qualitative Thesis fails. In any context, it is bad to say things like (34), repeated below:

(34) Either the plate won't be dropped or it will smash; but if it is dropped it might not smash.

This would give us:

> *High Standards QT.* If $\neg K \neg A$ and $K(A \supset C)$, then $K(A >_{HI} C)$

Given this principle, contextualism still leads to skepticism. The falsity of $A >_{HI} C$ suffices for me to fail to know it. And by *High Standards QT*, I also fail to know the material. The rest of the argument proceeds as before.

Similar things hold for the other two arguments. The second argument goes through given *High Standards QT* plus the principle that you know a future directed high-standards indicative only if you know the corresponding future directed high-standards counterfactual:

> *High Future Directed.* If at t you know a future directed indicative $A >_{HI} C$, then at t you know the corresponding future directed counterfactual $A \mathrel{\Box\!\!\rightarrow}_{HI} C$.

This also seems plausible. Again, there is *no* context in which it seems plausible to say things like (16), repeated below:

(16) #I know that the plate will smash, if it is dropped; but I don't know whether it would smash, if it were dropped.

Even given *High Standards CUW*, we now get a skeptical conclusion. By *High Standards CUW*, the counterfactual $A \mathrel{\Box\!\!\rightarrow}_{HI} C$ is false. So it is not known. By *High Future Directed* the indicative $A >_{HI} C$ is not known and the rest of the argument proceeds as in the previous paragraph.

The third argument also goes through given the high standards version of *Indicatives Turn Counterfactual*:

> *High Standards ITC.* If having total evidence E suffices for you to know $A >_{HI} C$, then having total evidence $E \cap \neg A$ suffices for you to know $A \mathrel{\Box\!\!\rightarrow}_{HI} C$.

Restricting ourselves to the kinds of cases where *Indicatives Turns Counterfactual* holds, *High Standards ITC* also seems plausible: in those contexts, the kinds of things that would undermine the counterfactual (chanciness, mightyness) undermine the indicative. Combining this with *High Standards CUW*, we can run the argument from §6 on the high standards counterfactual.

The common denominator in all three arguments is *High Standards QT*: if the contextualist denies this, they can block all three arguments. And from a theoretical point of view, the contextualist should find the principle to be strange. It's true that there is no context in which we can assert counterexamples to the Qualitative Thesis; but it seems that there *should* be counterexamples. To avoid skepticism, we must say we can know some proposition C, even when it has a small chance of being false; and, by Closure, often this will mean knowing certain disjunction $\neg A \vee C$. How could we know *both* $A >_{LO} C$ *and* $A >_{HI} C$ on the basis of $\neg A \vee C$? Our basis for knowing $\neg A \vee C$ is compatible with a small chance of it being false; but $A >_{HI} C$ cannot even be *true*, in such a case.

The natural solution is to appeal to the contextualist view of *knowledge* from the last section. As Ichikawa (2011, 2017) and Lewis (2017) have observed, the same sorts of considerations favouring counterfactual skepticism promote skepticism about knowledge also: if we make salient the small chance that the plate will not smash, we become less inclined to say that we know it will. This suggests the standards for "knows" are *coordinated* with the standards for counterfactuals and indicatives.[37]

Contextualism about knowledge can square this lack of counterexamples to the Qualitative Thesis with contextualism about conditionals. We might say that low standards knowledge of a disjunction goes with low standards knowledge of a low standards indicative; and high standards knowledge of a disjunction goes with high standards knowledge of a high standards indicative. This would give us two different versions of the Qualitative Thesis, where the standards are all coordinated:

Coordinated Low QT. If $\neg K_{LO}\neg A$ and $K_{LO}(A \supset C)$ then $K_{LO}(A >_{LO} C)$

Coordinated High QT. If $\neg K_{HI}\neg A$ and $K_{HI}(A \supset C)$ then $K_{HI}(A >_{HI} C)$

[37] This does not necessarily mean the relevant alternatives are the same for both; simply that changes the standards for one induces a corresponding change for the other.

Given these two principles, we have two versions of each of my three arguments, one about low standards conditionals and knowledge and another about high standards conditionals and knowledge. According to the contextualist, the first argument is valid, but not sound: the first premise fails, since low standards conditionals are true. The second argument *is* sound, but its conclusion should not trouble us: it merely shows that we do not have *high standards* knowledge of the future. But this is independently plausible: high standards knowledge of the future is inconsistent with a chance of falsehood.

Contextualism about conditionals now offers a response both to counterfactual skepticism and my arguments from earlier.[38] High standards conditionals are undermined by chance; low standards conditionals are not. So we fail to know high standards conditionals, but can know low standards conditionals. Given our two versions of the Qualitative Thesis above, this means we tend to *lack* high standards knowledge of the corresponding disjunction; but we will know it by low standards. By Closure, we cannot have high standards knowledge of anything that might entail such disjunctions. But that is unproblematic — contextualists about knowledge should already endorse this. We can, however, have low standards knowledge of the disjuncts; and that is all we ordinarily need.

My argument then shows us something interesting about contextualist responses to counterfactual skepticism. The contextualist thinks that the counterfactual skeptic is partly right: high-standards counterfactuals tend to be false, for the exact reasons the counterfactual skeptic gives. To avoid inheriting the problems of counterfactual skepticism, the conditional contextualist should also be a knowledge contextualist.

[38]Though not as satisfying as the epistemic response, to my mind at least. I am struck by the fact that there seems to be redundancy in the contextualist response above: it requires two forms of contextualism to do what the epistemic response achieves with one. (Moreover Matt Mandelkern (pc.) notes that something unusual happens here for the double-contextualist: changing the standards does not just render the knowledge claim false; it also renders the *prejacent* false.) Overall, I think the epistemic view does a more elegant job of reducing counterfactual skepticism to a general skeptical puzzle.

8 Conclusion

Our world turned out to be surprisingly chancy. Once that's accepted, asks the counterfactual skeptic, would it really be so bad if most counterfactuals were false? Perhaps this is the best way to reconcile our practice of using counterfactuals with a serious scientific world view: we say false things all the time; but we can always retreat to a probabilisitic counterfactual if challenged.

I have argued that the true picture is much worse than this. Given various plausible principles about conditionals, counterfactual skepticism is inconsistent with much of our knowledge of the future. Future skepticism is not part of any attractive world view, let alone a serious scientific one; and so I think counterfactual skepticism should be rejected. While there are genuine puzzles about we can have knowledge of the future, we also have the tools to make sense of those puzzles. My arguments give us independent reason to apply those tools to counterfactual skepticism too.

References

Adams, E. W. (1970). Subjunctive and indicative conditionals. *Foundations of Language*, 6:89–94.

Alchourrón, C. E., Gärdenfors, P., and Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530.

Bacon, A. (2015). Stalnaker's Thesis in context. *Review of Symbolic Logic*, 8:131–163.

Bennett, J. (2003). *Philosophical Guide to Conditionals*. Oxford University Press, Oxford.

Boylan, D. and Schultheis, G. (2022). The Qualitative Thesis. *The Journal of Philosophy*, 119:196–229.

DeRose, K. (1995). Solving the skeptical problem. *Philosophical Review*, 104:1–52.

DeRose, K. (2010). The conditionals of deliberation. *Mind*, 119:1–42.

Dudman, V. H. (1994). Against the indicative. *Australasian Journal of Philosophy*, 72:17–26.

Edgington, D. (1995). On conditionals. *Mind*, 104:235–329.

Edgington, D. (2008). Counterfactuals. *Proceedings of the Aristotelian Society*, 108:1–21.

Gibbard, A. (1981). Two recent theories of conditionals. In Harper, W., Stalnaker, R. C., and Pearce, G., editors, *Ifs*, pages 211–247. D. Reidel, Dordrecht.

Goldstein, S. and Santorio, P. (2021). Probability for epistemic modalities. *Philosophers' Imprint*, 21.

Grove, A. (1988). Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170.

Hájek, A. (2020). Contra counterfactism. *Synthese*, 199:181–210.

Hajek, A. (2022). Most counterfactuals are false. ms.

Hawthorne, J. (2003). *Knowledge and Lotteries*. Oxford University Press, Oxford.

Hawthorne, J. (2005). Chance and counterfactuals. *Philosophy and Phenomenological Research*, 70:396–405.

Hawthorne, J. and Lasonen-Aarnio, M. (2009). Knowledge and objective chance. In Greenough, P. and Pritchard, D., editors, *Williamson on Knowledge*, pages 92–108. Oxford University Press, Oxford.

Holguin, B. (2021). Indicative conditionals without iterative epistemology. *Noûs*, 55:560–580.

Ichikawa, J. (2011). Quantifiers, knowledge, and counterfactuals. *Philosophy and Phenomenological Research*, 82:287–313.

Ichikawa, J. (2017). *Contextualising Knowledge: Epistemology and Semantics*. Oxford University Press, Oxford.

Jackson, F. (1990). Classifying conditionals. *Analysis*, 50:134–147.

Khoo, J. (2022). *The Meaning of* If. Oxford University Press, New York.

Kocurek, A. W. (2022). Does chance undermine would? *Mind*, 131:747–785.

Lewis, D. K. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74:549–567.

Lewis, K. S. (2016). Elusive counterfactuals. *Noûs*, 50:286–313.

Lewis, K. S. (2017). Counterfactuals and knowledge. In Ichikawa, J., editor, *The Routledge Handbook of Epistemic Contextualism*, pages 411–424. Routledge, New York.

Loewenstein, Y. (2021a). Heim sequences and why most unqualified 'would'-counterfactuals are not true. *Australasian Journal of Philosophy*, 99:597–610.

Loewenstein, Y. (2021b). Should we be skeptics or contextualists about counterfactual conditionals? *Philosophy Compass*, 16.

Lycan, W. G. (2001). *Real Conditionals*. Oxford University Press, Oxford.

Mandelkern, M. (2018). Talking about worlds. *Philosophical Perspectives*, 32:298–325.

Moss, S. (2012). On the pragmatics of counterfactuals. *Noûs*, 46(3):561–586.

Moss, S. (2013). Subjunctive credences and semantic humility. *Philosophy and Phenomenological Research*, 87:251–278.

Nozick, R. (1981). *Philosophical Explanations*. Harvard University Press, Cambridge, MA.

Schulz, M. (2014). Counterfactuals and arbitrariness. *Mind*, 123:1021–1055.

Schulz, M. (2017). *Counterfactuals and Probability*. Oxford University Press, Oxford.

Stalnaker, R. C. (1975). Indicative conditionals. *Philosophia*, 5(3):269–286.

Stalnaker, R. C. (1981). A defense of Conditional Excluded Middle. In Harper, W., Stalnaker, R. C., and Pearce, G., editors, *Ifs*, pages 87–104. Reidel.

Steele, K. and Sandgren, A. (2020). Levelling counterfactual scepticism. *Synthese*, 199:927–947.

Stefánsson, H. (2018). Counterfactual skepticism and multidimensional semantics. *Erkenntnis*, 83:875–898.

Williamson, T. (2009). Probability and danger. In *Amherst Lecture in Philosophy*, pages 1–35.