

Should We Be Skeptics or Contextualists About Counterfactual Conditionals?

Forthcoming at Philosophy Compass

(This is the penultimate draft: please cite published version)

Abstract

Just as knowledge contextualism offers a way out of knowledge skepticism in the face of powerful skeptical arguments, counterfactual contextualism purports to answer the many compelling arguments for the skeptical thesis that most ordinary counterfactuals of the form ‘if A had happened, C would have happened’, are false. In this article I review a few of the arguments for counterfactual skepticism, before surveying the various types of contextualist responses. I then discuss some of the recent objections to counterfactual contextualism, with an eye toward weighing contextualism’s costs with the costs of accepting skepticism. I conclude by remarking on some of the implications that skepticism and contextualism each has for many disputes in metaphysics.

1. Introduction

What is the relationship between the truth-value of a counterfactual and the context in which it is uttered? The answer to this question is under active dispute. Over the last decade or so, there has been a building movement toward counterfactual contextualism: the view that, like with the paradigmatic context-sensitive expressions such as indexicals (*I, today*), and gradable adjectives (*tall, rich*), which propositions are expressed by counterfactual conditionals can vary according to the context in which they are uttered. The motivation for contextualism has largely come from two considerations: compelling arguments for counterfactual skepticism – the view that most ordinary ‘would’-counterfactuals of the form ‘if A had happened, then C would have happened’ are *false* – and the puzzle of Heim sequences. After briefly introducing counterfactuals and the role that context has traditionally been thought to play in their evaluation in §2, I will summarize a few of the arguments for counterfactual skepticism in §3. In §4 I will describe the Heim sequence puzzle. I will then survey the various types of contextualist accounts, how they reply to skepticism and the Heim sequence puzzle, and some of the problems with each, in §5. In §6 I will discuss some reasons to think that skepticism may not be as unpalatable as it may seem at first sight. I will conclude by remarking on some metaphysical implications of the skepticism-versus-contextualism debate in §7.

2. Counterfactuals and Context

It is widely agreed that context plays *some* role in the evaluation of counterfactual conditionals. Many semantic analyses, including the classic Lewis (1973), Stalnaker (1968) and Kratzer (2012) accounts, treat counterfactuals as context-sensitive quantifiers over possible worlds. On the Lewis-Stalnaker-Kratzer-style view, for instance, the accessible worlds in the domain of the counterfactual quantifier are ordered according to their similarity to the world of assessment, w_o , based on a particular, contextually-determined comparative similarity relation. If ' $\>$ ' represents the counterfactual connective, then on David Lewis's framework (but using my notation), for given propositions A and C, a counterfactual $A \> C$ is non-vacuously true just in case there is an (accessible) A & C-world (i.e., a world at which both A and C are true) that is more similar to w_o in the contextually-relevant respects than is any A & Not-C world. For instance, 'if I hadn't touched the bee with my finger, I wouldn't have been stung' is true just in case there is a possible world at which I don't touch the bee with my finger and I don't get stung which is *more similar* in the contextually-relevant respects to the actual world than is any possible world at which I don't touch the bee with my finger but I am still stung.

The role that context plays here is straightforward: it helps to determine the facts that are relevant to the world-similarity ordering. There are two considerations that have been especially influential in leading some theorists to think that context plays a much larger role than this. The first is the threat from counterfactual skepticism. The second is the puzzle of Heim Sequences. We'll take a brief look at each, in turn.

3. Skepticism

Counterfactual skepticism, a view made famous, and extensively defended, by Alan Hájek (Unpublished manuscript), (2020b), takes nearly all non-probabilistic 'would'-counterfactuals to be false. By non-probabilistic, I mean that there is no probability operator, or similar, in the consequent. (1P), below, is an example of a 'probabilistic' counterfactual:

(1) If Sophie had gone to the parade, she would have seen Pedro Martinez dance.

(1P) If Sophie had gone to the parade, she would have probably seen Pedro Martinez dance.

We can imagine that Sophie had been considering going to the parade at which Pedro Martinez was featured on a float, but that she had decided not to attend at the last minute. In that case both counterfactuals above seem acceptable (i.e., believable, assertible); both are of the sort that speakers use with regularity. It is counterfactuals like (1), though, that the skeptic thinks are false. That is because, unless the consequent is nomically or otherwise *entailed* by the antecedent, there will be seemingly-falsifying possibilities that cannot be ruled out. For instance, had Sophie gone to the parade, she might have been in the bathroom when Pedro's float went by. She might have been looking down at her phone at the time. She might have been stuck behind someone tall. And if that is right, then it is not clear how it can be true that if she had gone to the parade, she would have seen Pedro.

There are various ways to make the argument more precise. Given a similarity semantics, the idea can be put as follows.¹ Suppose the physical laws at w_o are indeterministic. In that case there is a world, call it *Nearest-No-See*, that is, up until Pedro Martinez's appearance at t_1 , *identical* in every way to *Nearest-See*: the most similar world to w_o at which Sophie goes to the parade and sees Pedro. Only at *Nearest-No-See*, she does not see Pedro (maybe a particular neuron doesn't fire and so she continues looking down at her phone). Plausibly, in that case, there is no feature of w_o that can make *Nearest-See* more similar to w_o than *Nearest-No-See* is – after all, the worlds are identical until t_1 , and if we are concerned about what happens at and after t_1 , then, if anything, *Nearest-No-See* is *more* similar to w_o , because like at w_o , Sophie doesn't see Pedro. If there is no world at which Sophie goes to the parade and sees Pedro which is more similar to w_o than is the most similar world at which she goes to the parade and does not see Pedro, then a similarity semantics will rule (1) *not* true.²

What if the relevant laws at w_o are deterministic? Then there will need to be some divergence between *Nearest-See* and *Nearest-No-See* prior to t_1 to account for the divergence at t_1 . As Nina Emery (2017: 405) points out, however, the differences might consist only in small changes in the exact positions of a group of relatively localized particles. Justification for thinking that the

¹ The following is adapted from Emery (2017).

² I.e., either *false* or *indeterminate*, depending on the particular semantics.

similarity function will be sensitive to these types of differences and that, even if it is, *Nearest-See* will be more similar than *Nearest-No-See*, is hard to come by.

In any case, we need not endorse a similarity semantics for the threat of skepticism to rear its head. Hájek (Unpublished manuscript) has a number of broadly-applicable arguments for counterfactual skepticism which don't take any particular semantic analyses for granted. One is the argument from "might"-counterfactuals. Because of possibilities like Sophie being in the bathroom at the time of Pedro's dance, the following seems undeniably true:

(1M[~]) If Sophie had gone to the parade, she might not have seen Pedro dance.

And yet, (1M[~]) clashes with (1):

(1M[~] & 1) # If Sophie had gone to the parade, she might not have seen Pedro dance; but she would have seen Pedro dance had she gone to the parade.

That sounds bad (the '#' indicates the infelicity of the utterance) and suggests that it is not the case that both are true. If that is right, and if (1M[~]) is undeniably true, (1) must be rejected.

We need not involve 'might'-counterfactuals to see the problem, either. Consider that for most counterfactuals in most contexts, the antecedent is under-specified: there are not enough details in the antecedent – even in conjunction with information from the context – for the event referred to by the consequent to determinately have occurred if the antecedent-event had. This is not an issue of indeterminacy in the *world* (though that, too, is the basis of another compelling argument for counterfactual skepticism (see Hájek (Unpublished manuscript))). The issue here is one of *semantic* indeterminacy. If Sophie had gone to the parade, at precisely what time would she have gone? At precisely what time would she have arrived, and exactly where would she have stood? When would she have seen something, or had some thought, that would have caused her to look down at her phone? It is highly implausible that all such facts can be fixed by information from the antecedent and context, alone. And if they are not fixed, then there is seemingly nothing to rule out that it is *not* the case that she would have seen Pedro for some reason or another.³

³ For more on these arguments, and for others in addition, see Hájek (Unpublished manuscript).

These arguments generalize. I have used (1) as the example, but it is as ordinary a counterfactual as any. I leave it as an exercise to the reader to apply the arguments to any other non-probabilistic ‘would’-counterfactual whose consequent is not entailed by its antecedent.

4. Heim Sequences

Before seeing how contextualists have responded to the skeptical challenge, we’ll take a quick look at the second motivation for counterfactual contextualism: the puzzle of Heim sequences.

Return to (1):

(1) If Sophie had gone to the parade, she would have seen Pedro dance.

Now consider what would have happened if Sophie had gone to the parade but had been stuck behind someone tall. (2) seems true as well:

(2) If Sophie had gone to the parade and been stuck behind someone tall, she would not have seen Pedro dance.

The classic semantics famously handles sequences like <(1) followed by (2)>, just fine. If all of the Sophie-goes-to-the-parade-worlds most similar to the actual world are worlds at which Sophie sees Pedro, then (1) is true.⁴ But (1)’s truth does not preclude (2)’s. (2) is also true if the most similar worlds at which Sophie goes to the parade *and* is stuck behind someone tall are worlds at which she does not see Pedro (we can put aside, for now, any skepticism over the idea that worlds at which Sophie is stuck behind someone tall are less similar to the actual world than worlds at which she sees Pedro).

But there is a problem. The problem of Heim Sequences, first discussed by Kai von Fintel (2001) and Anthony Gillies (2007), and attributed to Irene Heim, is that, assuming that there is nothing distinctive about Sophie to rule out her getting stuck behind someone tall (she’s not extraordinarily tall herself, for instance), if the order of assertion of (1) and (2) is reversed, then the second no longer rings true:

(2). If Sophie had gone to the parade and been stuck behind someone tall, she would not have seen Pedro dance.

⁴ For ease of exposition I speak here as though the ‘Limit Assumption’, the assumption that there is always a set of most similar antecedent-worlds, holds. Lewis (1973) denies the Limit Assumption but this is not important for our purposes.

(1). #But if Sophie had gone to the parade, she would have seen Pedro dance.

That (1) is infelicitous when asserted after (2) is unexpected on the standard semantics: on the standard picture it either is the case or it isn't that the most similar worlds at which Sophie goes to the parade are worlds at which she sees Pedro – the order of assertion has no role to play. So how do we explain the infelicity of (1) in this context?⁵

5. Counterfactual Contextualism

The contextualist's reply to both counterfactual skepticism and the puzzle of Heim sequences relies on the claim that which proposition a given counterfactual expresses is a context-sensitive matter. Just as *knowledge* contextualism offers a way out of knowledge-skepticism – the view that most ordinary knowledge-claims are false – by denying that purportedly-falsifying skeptical scenarios are relevant to the evaluation of knowledge-claims in ordinary contexts, counterfactual contextualists respond to counterfactual skepticism (and to the Heim sequence puzzle) by denying that certain counterfactual possibilities are relevant to the evaluation of counterfactuals in ordinary contexts. We can follow Hájek (2020b) in classifying counterfactual contextualist projects into three types: one can be a contextualist about the counterfactual connective, about the counterfactual's antecedent, or about its consequent.

a. Antecedent and Consequent-Contextualism

As one might expect, an antecedent-contextualist regards the proposition expressed by the antecedent of a counterfactual as context-sensitive. For instance, she may maintain that in most ordinary contexts – i.e., contexts in which very surprising or rare possibilities are not salient – (1) can be understood as expressing something like the following:

Naïve antecedent-contextualist reinterpretation of (1): If Sophie had gone to the parade *and nothing highly improbable or atypical had happened*, she would have seen Pedro.

Sandgren and Steele (2020) have developed a sophisticated antecedent-contextualist account. On their view, counterfactuals should be understood relative to a *domain of scientific inquiry*, although

⁵ Moss (2012) defends a *merely pragmatic* explanation for the infelicity which is consistent with (1) being true. Even if right, this would not help with the skeptical challenge, and a choice would still need to be made between skepticism and contextualism (barring another good alternative). For some objections to Moss's pragmatic solution to the Heim sequence puzzle see Lewis (2018).

this is to be interpreted loosely. Domains “mark out recognizable regularities or patterns in the world, be they of greater or lesser significance.” (2020:4). Some examples of domains include fundamental physics, human psychology, the dynamics of groups involving (a particular person named) Alice, and the promptness of service at a particular café. (2020:4) To interpret a counterfactual one identifies a domain of inquiry in the context. The counterfactual is then understood as follows:

Sandgren and Steele general interpretation: If A had happened *under fitting circumstances for domain d*, then C would have happened.

For example, suppose (3), below, is said about a reliably fun and personable person named *Alice*:

(3) Were Alice to come to the party, it would be fun.

The authors propose that in normal contexts the relevant domain of inquiry is ...plausibly one concerning Alice’s impact on others in ordinary social milieux. Here circumstances in which the party is ruined by a house fire or a plumbing problem are arguably not fitting, despite, perhaps, being relatively common events. (2020: 5)

In contrast to antecedent-contextualism, *consequent-contextualism* is the view that the proposition expressed by the *consequent* of a counterfactual is context-sensitive. Hannes Leitgeb (2012) has a well-developed version of this. A consequent-contextualist thinks that (1) means something like the following in ordinary contexts:

Naïve consequent-contextualist reinterpretation of (1): If Sophie had gone to the parade, she *probably/typically* would have seen Pedro.

Both antecedent- and consequent- contextualists offer a way out of skepticism. The purportedly falsifying possibilities – for instance, Sophie getting stuck behind someone tall and so not seeing Pedro, or Alice attending the party but there being a house fire – are now *consistent* with the relevant counterfactual, and so not falsifying after all.

The contextualists can also solve the Heim sequence puzzle. In typical contexts (1) is true. But if (2) comes first, it creates an atypical context by bringing to salience an improbable or unexpected event. In the new context, (1) is no longer to be interpreted with the contextually-modified antecedent or consequent – atypical possibilities have become relevant – and so (1) is false. On Sandgren and Steele’s view, the domain of inquiry has shifted to one which, among other things,

takes seriously the possibility of Sophie being stuck behind someone tall. Interpreted in the updated context, (1) is false.⁶

While both kinds of contextualist views can evade skepticism and account for Heim sequences, there are some problems. In my view, one of the most serious is simple: neither seems to accurately describe what speakers really mean when they use counterfactuals in ordinary contexts. There are a variety of ways to see that this is so. In a critical discussion of Leitgeb (2012), Hájek (2014) says the following:

Let us return to the putative near-synonymy of ‘would’ and ‘would very probably’. Suppose a doctor truthfully tells your father: ‘if you were to take this pill, you would very probably survive’; your father takes the pill and dies. Now suppose instead that the doctor had said: ‘if you were to take this pill, you would survive’; your father takes the pill and dies. Your case for a lawsuit is much stronger in the latter case than in the former. This also shows how Leitgeb’s semantics for counterfactuals allow modus ponens failures, which trouble me more than they appear to trouble him. ‘Would’ and ‘would very probably’ do not have almost the same meanings, because they license quite different inferences. (2014: 245)

The contextualist could reply that, given the high stakes, the doctor’s office is no ordinary context – however, the general worry that speakers don’t mean what the contextualists say they do can be illustrated using entirely ordinary counterfactuals in even the most benign of contexts. For instance, linguistic intuitions regarding retraction cases suggest that counterfactuals don’t typically mean what antecedent- or consequent- contextualists claim.⁷ Suppose that, unbeknownst to the conversational participants, Sophie actually attended the parade but was stuck behind someone tall and so did not see Pedro. Someone asserts (1): “If Sophie had gone to the parade, she would have seen Pedro”. When the truth is later discovered, the natural response is to think that the speaker had spoken falsely. We would expect the speaker himself to admit he had been wrong. This is surprising if the contextualists are right. As we’ve seen, on their view the truth of (1) is consistent with Sophie going to the parade but getting stuck behind someone tall (and so not seeing Pedro). If

⁶ Note that the skeptic also has an available solution to the Heim sequence puzzle: neither (1) nor (2) are technically true (although they are close to something that is true – see §6). However, (1) seemed true when it came before (2) in the sequence because the reasons that Sophie might have not seen Pedro did not occur to us. Once such reasons are made explicit through (2), we recognize (1)’s falsity.

⁷ The following is similar to an argument made in Loewenstein (2020). For a different kind of retraction argument against contextualism, see Hájek (2020b).

what the speaker had really meant when asserting (1) was, for instance, that if circumstances had been typical, Sophie would have seen Pedro, or that if Sophie had gone to the parade, she would have probably seen Pedro, then it should not seem that the speaker had spoken falsely. The fact that the assertion does seem clearly false is good evidence that the contextualists are wrong about (1)'s meaning in ordinary contexts.

Could *weak-centering* explain the falsity of (1) in a way which does not undermine the contextualists' interpretation? Weak-centering says that the world of assessment is always among the worlds closest to itself. The principle entails that if the antecedent of the counterfactual is true and the consequent false, then the counterfactual is false. If Sophie actually went to the parade but did not see Pedro, then weak centering seems to apply. The consequent-contextualist may indeed be able to appeal to weak-centering, though it will depend on the details of the account: in particular, it will depend on how the consequent-contextualist understands 'probable'. On an understanding of 'probable' according to which one can still say it is probable that Sophie would see Pedro were she to go to the parade, even if it turns out that she does go and she does not see him, weak centering *won't* be helpful to the consequent-contextualist. For in that case the consequent of the counterfactual (given the consequent-contextualist's understanding of it), is *true*: it *was probable* that Sophie would have seen Pedro had she gone to the parade. And if the consequent is true, weak-centering is inapplicable.

Retraction cases are even worse for the antecedent-contextualist, who, so far as I can tell, has no recourse available to her which will allow her to appeal to weak-centering. The problem for the antecedent-contextualist, in particular, is that if we interpret the antecedent as the antecedent-contextualist thinks we ought to, then if Sophie actually attends the parade but gets stuck behind someone tall, the proposition expressed by the antecedent *isn't* true: it *didn't* happen that Sophie went to the parade under fitting circumstances. So once again, weak-centering does not apply, and the apparent falsity of the assertion remains good evidence that the antecedent-contextualist is not giving an accurate account of what speakers mean.⁸

b. *Connective-Contextualism*

⁸ For some different objections to antecedent-contextualism, as well as some objections aimed specifically at Sandgren and Steele's account, see Hájek (2020b).

The connective-contextualist takes the function represented by the counterfactual connective, '>', to be context-sensitive. Karen Lewis (2016), (2018) has the best-developed account of this type.⁹ She proposes the inclusion of a contextual parameter into the semantics for counterfactuals that works in conjunction with the contextually-determined similarity function. The additional contextual parameter, the *relevance function*, is sensitive to features of the conversational context such as the standards of precision and the salience or non-salience of possibilities. It pushes worlds that are among the most similar but not relevant in the context out of the domain that the counterfactual quantifies over, and it also brings into the domain worlds that are not among the most similar but still relevant. Her truth conditions are below:

For all contexts c , $P \square \rightarrow Q$ is true at w in c iff all the closest P -worlds to w are Q -worlds, where closeness is a function of both similarity and relevance. (2018: 500)

This semantics evades the skeptical threat as follows. Although conversational participants might be concerned with fluky occurrences when conversing in a philosophy classroom, for instance, in ordinary contexts possibilities like Sophie getting stuck behind someone tall are typically not relevant. The relevance function pulls these worlds out of the domain of closest worlds, and (1) comes out true.

This is not to say that unusual or unexpected possibilities cannot be made relevant in a context. They can be made relevant, for example, by being brought to salience, which is exactly what is going on with Heim sequences. When someone utters (2) before (1) is asserted, it brings the possibility of Sophie getting stuck behind someone tall into the *common ground* – the information that conversational participants take for granted as background information for purposes of the conversation. Now relevant, the relevance function does *not* pull the worlds at which Sophie is stuck behind someone tall out of the domain of closest worlds, and (1) becomes false. Although (1) is true when asserted in the context prior to (2), it is false when asserted after it.

Connective-contextualism is arguably the most radical of the bunch. It is the version that takes counterfactuals to be most similar to the paradigm context-sensitive expressions like the gradable adjectives (*tall*, *wealthy*), since it takes the function itself – the function, in this case, from the propositions expressed by the antecedent and the consequent to the proposition expressed by

⁹ See also Ichikawa (2011), for an alternative connective-contextualist account.

the counterfactual – to vary based on the standards at play in the context. But are counterfactuals really like the context-sensitive expressions in the relevant respects? While it is clear that the truth-value of a statement like <Gertrude is tall> or <Domingo is wealthy> depends on the standards for tallness, or for wealthiness, in the context; intuitively it is *not* the case that whether Sophie would have seen Pedro dance had she gone to the parade depends on standards of any kind. Once the similarity ordering has been established, whether Sophie would have seen Pedro seems to have nothing to do with what occurs at the context of utterance. Sandgren and Steele (2020) put the point this way:

If K. Lewis is right, facts about counterfactual dependence seem to rock and sway with what groups of humans are interested in talking about. One might worry that the set of facts about what counterfactually depends on what is not so changeable and tying counterfactual dependence to conversational purposes compromises the status of counterfactuals as genuinely worldly statements about the way things are...On the other hand, we propose that the facts about counterfactual dependence are fixed by facts about the regularities in the world. (2020:19)

And Loewenstein (2020) emphasizes the difference between counterfactuals and paradigm context-sensitive expressions by contrasting discourses involving both:

[(4a)] He is tall!

[(4b)] [#?] No, he is not. Remember the NBA players whom we saw at the game?

[(4c)] _{ok} Don't be a smart aleck! I meant that he is tall for a boy his age.

[(5a)] The table is flat.

[(5b)] [#?] No, it is not; if you look with a microscope, you'll see unevenness.

[(5c)] _{ok} What I meant was that, for our purposes, it is flat. It is flat enough to sit your drink on it.

Compare with:

[(6a)] If Sophie had gone to the parade, she would have seen Pedro.

[(6b)] _{ok} But if Sophie had gone to the parade and had been stuck behind someone tall, she wouldn't have seen Pedro.

[(6c)] # That's not what I meant. I meant that, for our purposes, she would have seen Pedro./# Okay, relative to my standards she would have seen Pedro but relative to your standards she might not have./# That's not what I meant, I meant that if Sophie had gone to the parade she would have seen Pedro relative to standard (or restricted by domain) _____. [where the '____' is instantiated by some standard (or domain)]. (2020: 6)

The problematic replies in (6c) are evidence that the speaker who asserts (6a) takes herself to be speaking in absolute terms. She does not intend to assert that <if Sophie had gone to the parade, she would have seen Pedro> is true relative to some standard or other.¹⁰

6. Taking Skepticism Seriously?

We've now seen a few of the objections to the various contextualist projects. These and others should be weighed against the costs of accepting counterfactual skepticism. The main cost of the latter is that it violates a plausible principle of charity: we ought to interpret speakers as speaking truthfully when possible. Are we really willing to accept that most of what competent speakers say counterfactually is *false*? On the other hand, is attributing widespread falsehood really so bad if what speakers say is *close enough* to something that is true, and if, furthermore, communication purposes are served satisfactorily? On Hájek's (Unpublished manuscript) view, although most non-probabilistic 'would'-counterfactuals are false, there are nearby probabilistic counterfactuals of the form 'A-would-probably-C' that are true. While 'if Sophie had gone to the parade she would have seen Pedro' is false, 'if Sophie had gone to the parade it is very likely that she would have seen Pedro', is true. So, although speakers technically speak falsely when asserting most non-probabilistic 'would'-counterfactuals, they are asserting near-truths.

Indeed, there are reasons to think that using non-probabilistic 'would'-counterfactuals regularly is exactly what speakers *would* do even if Hájek is right about such counterfactuals being false. And if that's the case, then the available evidence – the fact that speakers regularly use non-probabilistic 'would'-counterfactuals – does not tell against skepticism. What are some of these reasons? As Elgin (2004), Braun and Sider (2007), and Teller (2011) have shown, imprecision is ubiquitous in ordinary language. People are rarely concerned about speaking as accurately as possible. What's more important, perhaps, is that it is *particularly* easy and common to drop terms like "probably". I might say "I am going to the park tomorrow" rather than "I am probably going to the park tomorrow", even though if pressed I will not hesitate to agree that the action is merely probable (the weather could become unexpectedly bad, I could come down with something, I could change my mind).¹¹

¹⁰ For additional objections to Lewis's contextualism see Hájek (2020b) and Loewenstein (2020).

¹¹ Of course, the speaker who drops the 'probably' here is not thereby saying something *false*, whereas the skeptic thinks that the person who drops the 'probably' in the consequent of the

We fail to weaken our counterfactual assertions, Loewenstein (2020) suggests, for just the reasons that we often fail to weaken our non-counterfactual ones:

If I were to get into the habit of always adding ‘probably’ before any counterfactual consequent, I would...inadvertently communicate much that I do not mean. If I were to say, “if that match had been struck, it probably would have lit”, my listeners might assume that I had some particular reason to think that the match might not have lit. Maybe I know something about the match, or its environment, that they do not. My listeners are likely to assume that I take the odds that the match would have lit...to be lower than I actually take them to be. To weaken one’s counterfactual assertion in this way could suggest some degree of uncertainty that might not be felt. For these reasons, and because it can be tiresome and unnecessary for speakers to be precise, and because the unlikely, unexpected ways one’s utterance could be made false are often not at the forefront of one’s mind when one speaks, it is not at all surprising that ‘A>C’ is regularly said in place of ‘A>probably-C’, even if the former is merely *close* to something that is true. (2020: 12-13)

7. Skepticism, Contextualism and Metaphysics

Whether skepticism or contextualism is the better choice, it is worth considering some of the implications that each has for other philosophical topics for which counterfactuals play an important role. There has been relatively little said on this, though Emery (2017) is a notable exception. She shows that the mere threat of skepticism has the surprising consequence of what she calls *chance pluralism*: the view that there may be multiple different chance values that attach to one given event at one world at one time. (2017: 400). The idea is that, since counterfactuals play an essential role in scientific inquiry, philosophical reasoning and ordinary discourse, either a way needs to be found to get most seemingly-true ‘would’-counterfactuals to come out true despite the skeptical arguments, or a way for most seemingly ‘would’-counterfactuals to still be *assertible*, even if untrue. And it turns out that finding a way – any plausible way – for counterfactuals to be either true or assertible in the face of the skeptical arguments, leads right into chance pluralism.

While this particular conclusion may seem surprising, it is *not* surprising that skepticism and contextualism both have important implications for a wide range of philosophical issues, especially in metaphysics. After all, counterfactuals are closely tied to many philosophical concepts, including causation, the laws of nature, moral responsibility, and many others. I won’t be able to do much by

counterfactual, is. But this difference is not a problem for the argument that speakers drop the ‘probably’ for the same reasons in both cases if speakers don’t *realize* that there is this difference. And surely most don’t.

way of investigating implications here, so will just make a couple of remarks. First, depending on the details, counterfactual skepticism may not have as big of an impact on metaphysical issues as one might think. For concepts for which we wish to give counterfactual analyses (or for which ‘would’-counterfactuals otherwise play an essential role), we will need to replace analyses-counterfactuals of the form ‘A-would-C’ with probabilistic counterfactuals of (something like) the form ‘A-would-(sufficiently)-likely-C’, since it is the latter that are true. For instance, instead of analyzing causal dependence between actual events A and C in terms of counterfactuals of the form ‘If A hadn’t happened, C wouldn’t have happened’, counterfactual analyses of causation will need to analyze causal dependence in terms of counterfactuals like, “If A hadn’t happened, it is (sufficiently) likely that C wouldn’t have happened.” This is a less drastic change than it may seem, given that those favoring counterfactual analyses of various kinds – including counterfactual analyses of causation – already face pressure to accommodate chancy causation, and usually the accommodations take just this form.¹²

The implications of contextualism are less straightforward. Would a commitment to counterfactual contextualism require a commitment to contextualism about anything for which one holds a counterfactual-based view? Consider causation again. Most take causation and counterfactuals to be essentially tied in some way. Would a commitment to counterfactual contextualism commit most causal theorists to causal contextualism, too? And, since causal responsibility is a necessary condition for moral responsibility, does a commitment to causal contextualism of the relevant sort commit one to the view that whether an agent is morally responsible can depend on what is salient in the context in which the moral responsibility assessment is made? Can we make it harder to be morally responsible by bringing atypical events into salience in the context of assessment, so that the relevant counterfactual claims, and, in turn, the relevant causal claims (and thus the relevant responsibility claims), come out false? There are different ways the contextualist could respond. She could answer “yes” to these questions. Alternatively, she could

¹² For example, in response to the possibility of chancy causation, Lewis amends his famous counterfactual analysis of causal dependence, originally: “Where *c* and *e* are distinct actual events, *e* causally depends on *c* if and only if, if *c* were not to occur *e* would not occur” (1973), so that instead it reads, “Where *c* and *e* are distinct actual events, *e* causally depends on *c* if and only if, if *c* had not occurred, the chance of *e*’s occurring would be much less than its actual chance” (1986).

maintain that contexts in which moral responsibility (for instance) is being evaluated are always special contexts with their own rules for what is and is not relevant to the evaluation of the counterfactual at issue. What is clear is that, given the important role that counterfactuals play in many philosophical disputes, some kind of story is owed for how contextualism about counterfactuals will impact other issues in metaphysics and philosophy more generally.¹³ Whether or not an acceptable story can be given may itself be important data for the choice between skepticism and contextualism.

Work Cited

- Braun, D. and Sider, T. (2007). "Vague, so Untrue", *Noûs* 41(2): 133 – 56.
- Elgin, C. Z. (2004). "True Enough", *Philosophical Issues*, 14: 113-21.
- Emery, N. (2017). The metaphysical consequences of counterfactual skepticism. *Philosophy and Phenomenological Research*, 94(2), 399-432.
- Gillies, A. (2007). "Counterfactual Scorekeeping", *Linguistics and Philosophy* 30(3): 329-60.
- Hájek, A. (2014). "Probabilities of Counterfactuals and Counterfactual Probabilities", *Journal of Applied Logic* 12 (3), 235-251.
- Hájek, A. (Unpublished manuscript). *Most Counterfactuals are False*. Accessed 2019.
- Hájek, A. (2020b). "Counterfactual Skepticism and Antecedent-Contextualism", *Synthese*, <https://doi.org/10.1007/s11229-020-02686-0>.
- Ichikawa, J. (2011). "Quantifiers, Knowledge and Counterfactuals", *Philosophy and Phenomenological Research*, 82(2): 287-313.
- Kratzer, A. (2012). *Modals and conditionals: New and revised perspectives* (Vol. 36). Oxford: Oxford University Press.
- Leitgeb, H. (2012). "A Probabilistic Semantics for Counterfactuals", *The Review of Symbolic Logic*, 5(1): 26-121.
- Lewis, D. K. (1973a). *Counterfactuals*. Oxford: Blackwell.
- Lewis, D. K. (1973b). "Causation", *Journal of Philosophy*, 70: 556-67.
- Lewis, D. K. (1986). "Postscripts to Causation" in his *Philosophical Papers Volume II*. Oxford: Oxford University Press.
- Lewis, K. S. (2016). Elusive Counterfactuals, *Noûs* 50(2): 286-313.
- Lewis, K. S. (2018). Counterfactual Discourse in Context, *Noûs* 52(3), 481-507.
- Loewenstein, Y. (2020) "Heim Sequences and Why Most Counterfactuals are Not True", *Australasian Journal of Philosophy*, DOI: 10.1080/00048402.2020.1792954.

¹³ Karen Lewis (2016: 310-311) makes some brief remarks on how such a story might go. More needs to be said, however, if we are to be in position to evaluate the account's plausibility in this respect.

- Moss, S. (2012), "On the Pragmatics of Counterfactuals", *Noûs* 46(3), 561-586.
- Sandgren, A. and Steele, K. (2020). "Levelling Counterfactual Scepticism", *Synthese*,
<https://doi.org/10.1007/s11229-020-02742-9>.
- Stalnaker, R. C. (1968). A theory of conditionals. *Studies in logical theory* (Vol. 2, pp. 98-112)., American philosophical quarterly monograph series. Oxford: Blackwell.
- Teller, P. (2011). "Two Models of Truth", *Analysis* 71(3): 465-72.
- von Fintel, K. (2001). "Counterfactuals in a Dynamic Context", in *Ken Hale: A Life in Language*, ed. Michael Kenstowicz, Cambridge, MA: The MIT Press: 123-52.