

Knowledge in the Face of Conspiracy Conditionals

Ben Holguín

April, 2018

1 Introduction

As I will understand the view, moderate invariantism about knowledge is the combination of two theses: (i) the moderate part: that we are generally in a position to know what is believed on the basis of perception, memory, and testimony; and (ii) the invariantist part: that if $\lceil S$ knows that $p \rceil$ is true in one context, it is true in every context.^{1,2} For the purposes of this paper we may abstract away from discussion of the moderate invariantist's preferred theory of the conditions under which perception, memory, and testimony suffice for knowledge. What matters are the specific knowledge *attributions* the moderate invariantist wants to deliver—i.e., the kinds of case by case judgments her theory is designed to predict. And on that matter I take it as uncontroversial that for an adult S in reasonably normal circumstances, the following are all paradigm cases of propositions of which S has (moderately invariant) knowledge:³

(P₁) That S has hands.

(P₂) That Donald Trump is the president of the United States.

(P₃) That Oswald shot Kennedy.

The core argument of this paper is that if a certain independently plausible principle about the felicitous use of indicative conditionals is true, then there is at least one context

¹ Given the potential context-sensitivity of the expressions S and p , this condition isn't quite right. More accurate would be something like: If $\lceil S$ knows that $p \rceil$ is true in one context c , then it is true in every context c' that assigns the same semantic value to S and p as c . But I will stick with the simpler construal in the main text for the sake of readability.

² What kinds of things are contexts? We need not have a settled answer. I am happy to follow in the spirit of (Kaplan, 1989) and model contexts as sequences consisting of something like: a speaker, a time, a modal base, a variable assignment, and, if need be, an epistemic index to allow for variation in the semantic contribution of epistemic vocabulary. What matters for the purposes of this paper is that contexts are *not* partially individuated in terms of (e.g.) the underlying epistemic facts. So in saying that $\lceil S$ knows that $p \rceil$ is true in c but false in c' , we can help ourselves to the assumption that the situations or circumstances or whatever you want to call them are constant across c and c' with regards to S 's beliefs and evidence, as well as the truth of matters relevant to p .

³ The point would be cleaner but more cumbersome to state if (P₁)-(P₃) were made tenseless.

c (in fact many) in which for each of the above propositions (P₁)-(P₃), the proposition expressed by $\lceil S$ knows that (P _{n}) \rceil is false in c . By extension, then, if the principle is true, then moderate invariantism about knowledge is false.

What's the principle? Put roughly, it's that there is something generally strange about asserting an indicative conditional \lceil If p , q \rceil when you know whether p . Thus, since conditionals like (1)

- (1) If Oswald didn't shoot Kennedy, then someone else did.

seem unproblematically assertable in a variety of ordinary contexts—or so I will argue—then contra moderate invariantism, it appears that there are ordinary contexts in which speakers don't know whether Oswald shot Kennedy.

It will then be argued that given the considerations motivating the principle, the best explanation of the existence of contexts in which it appears that (e.g.) $\lceil S$ knows that Oswald shot Kennedy \rceil is false is that 'knows' is context-sensitive. Views that give an alternative diagnosis of these facts—say that the appearance of (1)'s felicity is misleading, or that it is part of the semantics of 'knows' that those who are inclined to find (1) felicitous context-invariantly *fail* to know that Oswald shot Kennedy—are not plausible. Thus, the best account of our intuitive judgments about the felicitous use of indicative conditionals requires a contextualist theory of knowledge.

Here is the plan for the paper. §2 defends at length the principle that connects our judgments about the acceptability of indicative conditionals to facts about what is known in context. §3 gives some reasons to restrict the scope of the principle to so-called "canonical" uses of the indicative conditional, and explains what the distinction amounts to. §4 then presents a battery of examples of (what I will call) *conspiracy conditionals*—(1) being a paradigm case of such a conditional—and argues that they combine with the principle about indicatives to show that moderate invariantism about knowledge is false. §5 considers and rejects the view that the conspiracy conditionals of §4 can be excluded along the lines discussed in §3. §6 considers some analogs to the core principle that do *not* appeal to facts about what is known to generate the relevant judgments about the felicity of indicative conditionals. §7 examines the moderate invariantist's last line of defense: an error theory about the felicity judgments of conspiracy conditionals. It argues that such an error theory is implausible in a way the moderate invariantist's more familiar error theories are not. Finally, §8 argues that the problems raised by conspiracy conditionals are just as worrisome for proponents of *skeptical* and *subject-sensitive* invariantism about knowledge. So, the paper concludes, a proper account of our intuitions about indicative conditionals requires the resources of epistemic contextualism.

2 IGNORANCE

We begin with IGNORANCE, the principle that connects our judgments about the felicitous uses of the indicative conditional to facts about what is known. Put roughly, the principle is that whenever one knows whether p , there is in general something strange about asserting an indicative conditional whose antecedent is p . Why ‘in general’ rather than, say, ‘always’? Because as we will discuss in detail in §3, there is a distinction between *canonical* and *non-canonical* uses of the indicative conditional, and the weirdness typically only arises for canonical uses. But more on that later.

With that caveat in the background, here is a precise statement of the core principle:

IGNORANCE For canonical uses of the indicative ‘If p , q ’: In every context c : if ‘ S knows whether p ’ is true in c , then ‘ S may *not* felicitously assert the indicative conditional ‘If p , q ’ is true in c .⁴

Here I understand ‘felicitously’ in such a way that an assertion is felicitous iff the content of the assertion is neither false nor presuppositionally defective, and the assertion itself does not carry problematic implicatures.^{5,6}

The motivating thought behind IGNORANCE is that the primary purpose of an indicative conditional is for reasoning and talking about uncertainties: propositions whose truth is not settled by what you know. If you’ve settled whether p , you shouldn’t be asserting things like ‘If p , q ’. Instead, you either should be asserting q (when you know that p), or you should be asserting the counterfactual ‘If it were that p , it would be that q ’ (when you know that $\neg p$). That is the case for IGNORANCE stripped to its core. More can be said in the principle’s favor—and that is what the rest of this section will do—but the arguments are in some sense all variations on the same theme.

Before getting to the more detailed arguments, it will be helpful to have on the table some principles about the relationship between knowledge and assertion. Here are three closely related ones that will appear at various points in later discussion:

KSA In every context c : If ‘ S knows that p ’ is true in c , then ‘ S may (epistemically) assert p ’ is true in c .

⁴The principle is stated meta-linguistically so that contextualists and invariantists alike can make use of it. But I leave open that one’s preferred theory of knowledge or the felicitous use of indicative conditionals could have it that the appeal to contexts is superfluous.

⁵This is a technical use of ‘felicitous’. On the way I am understanding the notion, it is possible for an utterance to “sound fine” even when it is (strictly speaking) infelicitous. I will be careful to distinguish an utterance’s *being* felicitous from *seeming* felicitous when the difference might matter.

⁶Notice that there is no requirement that the conditional be *true*. It just can’t be false. The paper’s core argument is thus consistent with non-factualism (alternatively: non-propositionalism) about indicative conditionals. See, e.g. (Edgington, 1995) for a paradigm non-factualist theory of indicatives, as well as (Rothschild, 2012) for a helpful overview of the issues involved here.

KNA In every context c : If $\lceil S$ may (epistemically) assert $p \rceil$ is true in c , then $\lceil S$ knows that $p \rceil$ is true in c .

—and their conjunction:

KA In every context c : $\lceil S$ may (epistemically) assert $p \rceil$ is true in c iff $\lceil S$ knows that $p \rceil$ is true in c .⁷

In words: **KSA** says that knowledge is (epistemically) *sufficient* for assertion; **KNA** says that knowledge is (epistemically) *necessary* for assertion; and **KA** is the conjunction of **KSA** and **KNA**. The ‘epistemically’ qualifier is meant to control for cases in which an assertion of p would be problematic for reasons that go beyond the speaker’s evidential grounds for p . These are cases in which the assertion of p is impolite or otherwise harmful, as well as cases in which the assertion is pragmatically problematic—say because it is too weak, irrelevant, misleading, etc. I will not defend any of these principles about knowledge and assertion here, though I will make some suggestive remarks in their favor later in the paper (§6).⁸ Some of the arguments for **IGNORANCE** will make use of **KSA**, others will make use of **KNA**, and some will make use of neither. But having the principles on hand will make all of this easier to track.

The main argument in favor of **IGNORANCE** is that conditionals with antecedents whose truth values are known tend to sound somewhere between strange and terrible. Here is an example. Suppose your (otherwise normal) friend Jane comes up to you and says:

(2) ? If I had a bagel for breakfast this morning, then I went out for brunch with Jim.

If you are like me, you will find such an utterance puzzling. The thoughts that suggest themselves are: (i) that Jane has forgotten what she ate for breakfast; or (ii) that Jane is trying to engage you in some kind of unpleasant puzzle about her daily activities. The same is true of:

(3) ? If I didn’t have a bagel for breakfast this morning, then I ate cereal at home.

Supposing you know Jane isn’t the kind of person to speak in roundabout ways about what she ate for breakfast—i.e., that (ii) isn’t happening—you might be inclined to respond along the lines of:

(4) What do you mean *if* you had [didn’t have] a bagel for breakfast this morning? Do you not remember what you ate for breakfast?

⁷ The formulation of **KA** (including the qualifier ‘epistemically’) is inspired by [Worsnip \(2017, p. 6\)](#).

⁸ For some arguments in favor of there being a connection between knowledge and assertion along the lines of **KNA** (though not necessarily in exactly this form), see (e.g.) [DeRose \(2002\)](#); [Unger \(1975, ch. 6\)](#); [Williamson \(2000, ch. 11\)](#). See also ([Worsnip, 2017](#)) and the citations therein for a sense of the larger debate around knowledge accounts of assertion, with particular attention to its relationship to epistemic contextualism.

The fact that (2) and (3) seem strange and that (4) seems like an appropriate response admits of a straightforward explanation: when someone utters an indicative conditional $\lceil \text{If } p, q \rceil$, it is appropriate to infer that that person does not know whether p . Since you would expect Jane to know whether she ate a bagel for breakfast, her utterance of an indicative like (2) or (3) is jarring. This is what IGNORANCE is supposed to reflect.⁹

It is also worth noting how much the conditionals improve if the breakfast under discussion isn't this morning's, but one from a few weeks ago:

- (2*) If I had a bagel for breakfast two Saturdays ago, then I went out for brunch with Jim.
- (3*) If I didn't have a bagel for breakfast two Saturdays ago, then I ate cereal at home.

Why the improvement? Because obviously there is no general presumption that people know what they ate for breakfast weeks ago. IGNORANCE correctly has no complaint about an assertion of either (2*) or (3*).

A related bit of evidence for IGNORANCE comes from the oddity of speeches of the form $\lceil p, \text{ but if } \neg p, q \rceil$:

- (5) ?? Lexie is playing viola right now. But if Lexie is playing guitar right now, she isn't playing viola.
- (6) ?? It rained every day last week. But if it was sunny on Wednesday, then it was probably sunny on Thursday too.

I'm not sure whether these speeches ought to be classified as incoherent or merely very strange, but I'm going to take it as a datum that there is something definitely not right about them.¹⁰ Given KNA—the principle that says knowledge is (epistemically) necessary for assertion—IGNORANCE provides a simple explanation of the felt oddity of (5)-(6). And that is: (i) if you are in a position to assert p , then you will be taken to know that p ; (ii) given IGNORANCE, if you are in a position to assert $\lceil \text{If } \neg p, q \rceil$, then you will be taken *not* to know that p ; thus (iii), in asserting $\lceil p \text{ but if } \neg p, q \rceil$ you represent yourself as both knowing p and not knowing p . Hence why (5)-(6) seem bizarre.

A final empirical consideration comes from the following sorts of question/answer pairs:

- (7) a. Might Jake be at home? (Is it possible that Jake is at home?)
b. Yeah, if he's sick he's at home.
- (8) a. Must Jake be at work? (Is it certain that Jake is at work?)
b. No, if he's sick he's at home.

⁹ Notice that this explanation invoked neither KSA nor KNA.

¹⁰ Though see the discussion of so-called 'echoing' and 'concessive' uses of the indicative conditional in §3.

It is difficult to see why (7b) would constitute an appropriate affirmative answer to (7a) if not for the fact that an assertion of an indicative conditional is a way of conveying information about what might be the case. Likewise, it is difficult to see why (8b) would constitute an appropriate negative answer to (8a), if not for the fact that a context in which ‘If he’s sick, he’s at home’ is assertable is one in which it’s not the case that Jake must be at work. The natural conclusion to draw from these data is that epistemic modals and indicative conditionals are assessed against the same kinds of possibilities. Given the (relatively) uncontroversial assumption that the possibilities relevant in the assessment of epistemic modals are in most cases determined (partly) by what is *epistemically* possible for the speaker—i.e. what is compatible with what the speaker knows—we get a clear connection between IGNORANCE and our judgments about (7)-(8).¹¹ An affirmative answer to the question ‘Might Jake be at home?’ is explained by an utterance of ‘If he’s sick, he’s at home’ precisely because the conditional is assertable only if it is an epistemic possibility that Jake is sick. If it’s an epistemic possibility that Jake is sick, and the conditional ‘If he’s sick he’s at home’ is assertable, then it must also be an epistemic possibility that Jake is home. Thus, Jake might be at home. The explanation for why ‘If he’s sick, he’s at home’ constitutes a negative answer to ‘Must Jake be at work?’ is essentially the same. ‘If he’s sick, he’s at home’ is assertable only if it’s an epistemic possibility that Jake is sick. If it’s an epistemic possibility that Jake is sick, and if ‘If Jake is sick, then he’s at home’ is assertable, then it’s an epistemic possibility that Jake is at home. Hence it’s not the case that Jake must be at work.

Lastly, it is worth making a general point about why we should expect IGNORANCE to be true, at least when we restrict attention to the sorts of conversational settings those working in the tradition of, e.g., Grice (1975) and Gazdar (1979) take themselves to be theorizing about. Here the considerations separate somewhat neatly between cases of indicatives with antecedent that are known to be true and cases of indicatives with antecedents that are known to be false.

Starting with the first case, suppose you know that p and are considering whether to assert ‘If p , q ’. There are three possibilities: either you know that q , you know that $\neg q$, or you know neither. No matter which obtains, there will be something peculiar about an assertion of ‘If p , q ’. If you know that q , then an assertion of ‘If p , q ’ is problematically weak. Why assert the conditional when you could have just asserted q (or even ‘ p and q ’)? It would be strange much in the way it is strange to say ‘Some of my students failed’ when you know that all of your students failed. And it is this sort of strangeness you pick up on when Jane asserts (2) rather than that she ate breakfast with Jim.¹² If instead you know

¹¹ I do not intend to be making any especially controversial assumptions about the semantics of epistemic modals here. The assumption is really quite weak: it’s just that at some point in the calculation of the semantic value of ‘Might/Must p ’, facts about what the speaker *knows* (rather than believes or entertains or what have you) are involved.

¹² A similar line of reasoning suggests that it should be generally impermissible to assert ‘ p or q ’ when either

that $\neg q$, then the problem with an assertion of $\lceil \text{If } p, q \rceil$ is even simpler: bracketing exotic putative counterexamples to modus ponens, the conditional is straightforwardly false, or is at least the kind of thing in which you should have approximately zero confidence. Finally, if you do not know whether q , then although you don't know whether $\lceil \text{If } p, q \rceil$ expresses a falsehood (or something in which you should have approximately zero confidence), you also don't know that it *doesn't*. And so by KNA it is unassertable.¹³

As regards to cases in which the antecedent of an indicative conditional is known to be false, the argument is more direct. If you know that $\neg p$ and want to assert a conditional whose antecedent is p , then rather than assert the *indicative* $\lceil \text{If } p, q \rceil$, you should assert the *counterfactual* $\lceil \text{If it were that } p, \text{ it would be that } q \rceil$.¹⁴ Support for this idea, at least in some form or another, is widespread in the literature on conditionals—see, e.g., Stalnaker (1975); von Fintel (1999); Gillies (2010); Leahy (2011); Khoo (2015). Indeed, many believe it is a presupposition of the indicative that its antecedent is an epistemic possibility. I will thus take this half of IGNORANCE to be relatively uncontroversial.¹⁵

That is the case for IGNORANCE. Setting aside the quality of the arguments just presented in its favor, the principle seems to me something of a platitude about the standard use of indicative conditionals. I have yet to encounter an alternative theory of the sorts of reactions we have to speeches like (2)-(8) in the literature, and I imagine that is because much of the work on conditionals takes something like IGNORANCE for granted. That there is a tension between moderate invariantism about knowledge and IGNORANCE is thus of interest in its own right, even for those who would give up the latter before the former.

disjunct is known, as well as $\lceil \text{If } p, q \rceil$ when the *consequent* q is known. I think these generalizations are basically as good as the one IGNORANCE is meant to capture, but the class of exception cases is messier for them than it is for IGNORANCE. Between that and the fact that IGNORANCE is more than enough to get the paper's central argument going, I will continue to focus on conditionals with *antecedents* whose truth value is known.

¹³ Strictly speaking nothing as strong as KNA is needed. All that is needed is (i) the principle that whenever you know that p but don't know whether q , $\lceil \text{If } p, \text{ might } q \rceil$ is assertable, and (ii) the principle that if $\lceil \text{If } p, \text{ might } q \rceil$ is assertable, then $\lceil \text{If } p, q \rceil$ isn't.

¹⁴ A note about terminology: I distinguish indicative conditionals from counterfactual conditionals wholly in terms of their *semantic* properties, rather than (e.g.) their syntactic properties. Putting things as roughly as our purposes call for: an indicative conditional is a conditional whose semantics concerns what is happening at a contextually determined set of *epistemic* possibilities, while a counterfactual conditional is a conditional whose semantics concerns what is happening at a contextually determined set of (sometimes epistemically impossible) *metaphysical* possibilities. Nothing in this distinction rules out the existence of conditionals with indicative marking (e.g., $\lceil \text{If } p, q \rceil$) whose semantics is counterfactual, and likewise conditionals with subjunctive marking (e.g., $\lceil \text{If it were/had been that } p, \text{ it would be/would have been that } q \rceil$) whose semantics is indicative. See, e.g., von Fintel (1999); Khoo (2015) and the citations therein for further discussion.

¹⁵ Though to be fair, the literature has not been especially concerned with offering an account of the epistemic possibilities involved in the alleged presupposition. Any such account must settle (at least) two issues. First, the question of the kind of *attitude* epistemic possibilities are in the business of tracking: is it knowledge, belief, etc.? And second, the question of the extent to which this attitude must be *shared* among conversational participants: is it common knowledge/belief, mutual knowledge/belief, merely the speaker's knowledge/belief, etc.? With regards to the first question, for now I am happy to follow many in the literature and simply assume that the relevant attitude is knowledge, rather than belief or some other placeholder notion. (§6 will argue for this assumption at length.) And with regards to the second question, not much changes depending on our answer. It is just as much a strike against moderate invariantism about knowledge that there are quotidian contexts in which it is not commonly known that various mundane propositions are true as it is that there are quotidian contexts in which an ordinary speaker fails to know it.

3 Non-canonical uses of the indicative conditional

Before giving the argument against moderate invariantism from IGNORANCE, it is important that we get clear on the nature of the qualifier about canonical uses of the indicative conditional. This section will present a brief taxonomy of non-canonical uses, and then say a bit about what unifies the conversational circumstances in which one tends to find them.

The most straightforward statement of the difference between canonical and non-canonical uses of the indicative conditional is this: when used *canonically*, an indicative conditional $\lceil \text{If } p, q \rceil$ is a device for reasoning about and expressing connections between propositions that are *open* in the relevant context. All other uses are non-canonical. What does it mean for a proposition to be open in a context? Well, that's a matter that can (and will in §6) be disputed, but the natural thought is that a proposition's being open is, unsurprisingly, a matter of neither it nor its negation being *known*. On this way of understanding the canonical/non-canonical distinction, canonical uses of $\lceil \text{If } p, q \rceil$ are those in which the speaker does not know whether p , while non-canonical uses are those when the speaker does know whether p .

Now obviously this way of characterizing the distinction isn't very helpful at the present stage of the dialectic. For one, absent some other means of grasping the canonical/non-canonical distinction, IGNORANCE threatens to express a triviality.¹⁶ For another, this characterization of the distinction will be dialectically fraught once the argument against moderate invariantism is given in §4. And that's because there we're going to want to use IGNORANCE plus our judgments about putatively canonical indicative conditionals to impose constraints on a theory of knowledge. One who doesn't like the resulting constraints could always go by way of modus tollens and argue the uses in question must not be canonical after all. So it will be useful to have some more impartial ways of applying the distinction.

With that concern in mind, I believe the easiest way to get a relatively theory-neutral grip on the distinction is through examples. Taking a rough stab at a reasonably comprehensive taxonomy, here are the (non-canonical) uses of the indicative conditional that I believe call for the distinction in question:

First there are **echoing uses**, e.g.:¹⁷

- (9) If Pete's report is accurate, then Peggy is working at home. But we know Peggy never works from home, so Pete's report must be inaccurate.

¹⁶ If a conditional's use is canonical only if its antecedent isn't known to be true or known to be false, then of course it's the case that if S knows whether p , then S 's use of $\lceil \text{If } p, q \rceil$ is non-canonical—and *ipso facto* non-canonical or infelicitous.

¹⁷ I get the 'echoing' label from Dorr & Hawthorne (2013, pp. 890-91). Note that Dorr and Hawthorne only discuss echoing uses of conditionals whose antecedents are known to be true, like (10), and not echoing uses of conditionals whose antecedents are known to be false, like (9). For further discussion of the latter kind of echoing use, they cite (Declerck & Reed, 2001).

(10) If Dave is here, Roger is at home. Dave is here, so Roger is at home.

Echoing uses tend to occur in circumstances in which the speaker is *rehearsing* (echoing) her evidence for various propositions she takes herself to have settled: in this case either the negation of the conditional's antecedent, as in (9), or the conditional's consequent, as in (10). More generally, echoing uses occur in what one might think of as "dialectical" contexts—situations in which one is trying to give an argument for a known conclusion, or in which one is trying to reason one's way from known premises to a (soon to be) known conclusion. Consequently, a simple test for identifying an echoing use of $\lceil \text{If } p, q \rceil$ in the wild is to ask whether in the same context a followup assertion of $\lceil p; \text{ so } q \rceil$ or $\lceil \neg q; \text{ so } \neg p \rceil$ would be appropriate; if and only if it would be are you dealing with an echoing use of $\lceil \text{If } p, q \rceil$.

Second there are **Dutchman uses**, e.g.:

(11) If Jeremy goes to the gym twenty hours a week, then I'm a Dutchman.

(12) If Emily can afford to quit her job, then I'm a monkey's uncle.

Dutchman uses are essentially just echoing uses with flare. They are used to communicate the negations of their antecedents.¹⁸

And third, there are **Concessive uses**, e.g.:

(13) I know I cannot afford this timeshare. But even if I'm wrong about that, I still don't want to buy it.

(14) I am not lying about where I live. And even if I am, you're not going to get a different answer from me.

Concessive uses of $\lceil \text{If } p, q \rceil$ tend to occur in circumstances in which the speaker has settled that $\neg p$ and q , but knows that it is unlikely that her audience will take her to know whether $\neg p$. In light of this the speaker must settle for a "backup" position, q . In these circumstances a speech like $\lceil \neg p; \text{ but even if } p, q \rceil$ can seem perfectly natural. Consequently, a simple test for identifying a concessive use of $\lceil \text{If } p, q \rceil$ in the wild is to ask whether the use is occurring in a conversational context in which the speaker can expect her interlocutors to take her to know the things she asserts; if and only if the speaker *can't* expect to be trusted to know whether p if she were to assert p are you likely dealing with a concessive use of $\lceil \text{If } p, q \rceil$.

Three quick observations about this taxonomy. First, as concerns the purposes of this paper it is basically comprehensive.¹⁹ Second, it should be clear that none of these uses

¹⁸ It is not difficult to imagine canonical uses of either (11) or (12). This will happen when one lacks knowledge of the falsity of the antecedent, but has especially high confidence in its falsity.

¹⁹ There are two other categories of conditionals whose uses are typically non-canonical, but whose discussion I omit from the main text given their failure to pose any clear problem for IGNORANCE: *biscuit conditionals* and *donkey conditionals*. Examples of biscuit conditionals include:

typically implicates that the speaker is ignorant of the truth value of the conditional's antecedent.²⁰ Third and most importantly, setting aside complaints about which uses of the conditional deserve the honorific 'canonical', there is an obvious difference between how the conditional is used in each of these cases, and how it is used in the natural contexts associated with (2)-(8). One doesn't need a theory of the difference to be able to see it, and it should not go unappreciated that these various non-canonical uses *have special labels*.

Those points aside, what seems to me to be in common to the various non-canonical uses of the indicative conditional is that they all arise in circumstances in which the usual pressure to assert the strongest propositions relevant to the topic of conversation has been superseded by other communicative pressures. In the circumstances in which echoing (and Dutchman) uses are natural, the primary concern is not with sharing one's views on the topic, but with making explicit one's evidence for one's views on the topic. And in the circumstances in which concessive uses are natural, again the primary concern is not with sharing one's views on the topic, but with sharing the views that one expects one's uncooperative interlocutors will actually accept. By contrast, when one is in the kind of conversation where speakers are expected to just *tell* one another their views about the topic under discussion, one needn't worry about rehearsing inferential connections or hedging.

The natural thought, then, is that canonical uses of the indicative conditional arise in conversational circumstances in which it is expected that the participants assert what they can about the topic under discussion in a manner that meets an appropriate balance between informativeness and brevity. As alluded to briefly in §2, I imagine that these are the sorts of conversational contexts theorists like Grice (1975) and Gazdar (1979) are trying to model the dynamics of. Whenever one is in a position to impart information to one's interlocutors and the question of whether *q* is relevant to the topic of conversation, one's

(15) There are biscuits on the sideboard if you want some.

(16) If you're interested James is playing at the Roadhouse tonight.

It should be clear that biscuit conditionals are unlike ordinary indicative conditionals, at least on the latter's standard uses. But regardless of whether or not they are canonical, they are not counterexamples to IGNORANCE. It is just as strange to assert a biscuit when one knows whether its antecedent is true as it is to assert an ordinary indicative when one knows whether its antecedent is true. (See, e.g., DeRose & Grandy (1999) and Predelli (2009) for further discussion of biscuits.) Examples of donkey conditionals include:

(17) If a farmer owns a donkey, he beats it.

(18) If a cat has a hat, it wears it.

Mysteries abound with donkey conditionals—see, e.g., (King & Lewis, 2017) for an introduction. I will not pay much attention to them here, as it is not obvious they even count as indicative conditionals in the sense I am interested in (see footnote 14 above). But even if they are genuine indicatives, then as far I can tell they are just echoing conditionals, and so what goes for those goes for these. But even if *that's* wrong, it should be clear that the conditionals of interest to this paper (to be introduced in §4) are relevantly dissimilar to conditionals like (17) and (18).

²⁰ In fact it is plausible that the opposite is implicated—i.e., that the speaker *knows* whether the antecedent is true.

assertion of $\lceil \text{If } p, q \rceil$ will implicate that one doesn't know whether p . If one knew that p , then one would have been doing more to further the purposes of the conversation to assert q ; if one knew that $\neg p$, then one would have used a counterfactual conditional. Having asserted neither q nor a counterfactual conditional, one can expect one's listeners to infer that one doesn't know whether p . Conversational contexts shaped by other goals are certainly possible (perhaps even quite common), but the uses of the indicative conditional they give rise to are nonetheless non-canonical.

I conclude that the distinction between canonical and non-canonical uses is both natural and recognizable enough to leave the theoretical interest of IGNORANCE unscathed. For what IGNORANCE tells us is that if an assertion of $\lceil \text{If } p, q \rceil$ is felicitous in a context, then either it occurred in conversational circumstances shaped by non-standard communicative pressures, or by the lights of the context it is false that the speaker knows whether p . And this means that so long as we can tell reliably enough what the communicative pressures of various conversational circumstances are, we can use IGNORANCE as a tool for determining the extension of 'knows' in a given context.

4 The argument against moderate invariantism

Having made the case for IGNORANCE in §2 and clarified the canonical/non-canonical distinction in §3, we can now turn to the argument against moderate invariantism about knowledge.

Recall that central to moderate invariantism about knowledge is the thesis that for a normal adult S in reasonably ordinary circumstances, $\lceil S$ knows that $P_n \rceil$ is true in *every* context when P_n is substituted for any of the following:

- (P₁) That S has hands.
- (P₂) That Donald Trump is the president of the United States.
- (P₃) That Oswald shot Kennedy.

With that in mind, let us say that an indicative conditional $\lceil \text{If } p, q \rceil$ is a *conspiracy conditional* just in case its antecedent expresses one of the various propositions (or the negations thereof) the moderate invariantist takes normal people to context-invariantly know.

The colorful label is due to the paradigm case of such a conditional, which, thanks to Adams (1970), also happens to be a (philosophical) paradigm case of *the indicative conditional*:

- (1) If Oswald didn't shoot Kennedy, then someone else did.

(1)'s antecedent expresses the negation of (P_3), and so by stipulation is a conspiracy conditional. It is not an exaggeration to say that (1) might be the most philosophically famous example of an indicative conditional.²¹ The reason for its fame is that it makes for an excellent case study in the semantic differences between indicative and counterfactual conditionals—for compare (1), which is intuitively true, to its counterfactual analog (19), which is intuitively false:

(19) # If Oswald hadn't shot Kennedy, someone else would have.

Somewhat surprisingly, however, (1)'s philosophical fame more or less ends here. The fact that its naturalness is in tension with a popular package of views on knowledge and the pragmatics of indicative conditionals has, to my knowledge, gone largely unnoticed.²² For again, *if* moderate invariantism and IGNORANCE are both true, then the only felicitous uses of (1) are non-canonical uses. We will get into the prospects of such a line in more detail in §5, but for now it suffices to say that (1) seems (on its most natural uses) both obviously felicitous and obviously canonical.²³

What makes (1) a conspiracy conditional is that its antecedent expresses a proposition the moderate invariantist takes us to know the truth of the matter of. In this case, we know that it is false. But it is just as easy to come up with a variant in which the antecedent is something the moderate invariantist takes us to know to be *true*, as in:

(20) If Oswald shot Kennedy, then Kennedy really did die on November 22nd, 1963.

By definition, (20) is as much a conspiracy conditional as (1). And like (1), a canonical use of (20) would be perfectly appropriate in a variety of ordinary circumstances.

Conspiracy conditionals are by no means limited to conspiracies. Here are two others that seem fine out of the blue:

(21) If Trump is still president of the United States, then he hasn't had a fatal heart attack in the past few moments.

(22) If Trump is no longer the president of the United States, then it is surprising we haven't heard about that.

Likewise, suppose you're driving through the country side. In at least some contexts it would be perfectly normal to express thoughts like:

(23) If in a few hours the car breaks down, then we won't make it to our destination tonight.

²¹ As of early 2018, it is the first example of an indicative conditional on the Stanford Encyclopedia of Philosophy's entry for "Indicative Conditionals" (Edgington (2014)).

²² The only exception I know of is Gillies (2004, p. 585), who mentions it once and then does not return to the issue.

²³ I take (1)'s paradigm status to be the decisive first word on this question, even if not the decisive last word. No one bats an eye at it when it is introduced in discussions of indicative conditionals.

- (24) If these barns are more than a few days old, then it is unlikely we are driving through a movie set.

Given the assumption that we're typically in a position to know where we'll be in a few hours (at least absent good reason to expect that something will go wrong) and that barns we pass by on our way through the country have been around for more than a few days, we get further violations of IGNORANCE from the conspiracy conditionals (23)-(24).

The same can be said for standard cases of knowledge through testimony. If you read in the newspaper that Federer defeated Murray in straight sets, then in at least some contexts it would be perfectly natural for you to express thoughts as:

- (25) If Federer defeated Murray in straight sets, then this report is accurate.
(26) If Federer didn't defeat Murray in straight sets, then this report is inaccurate.

Finally, philosophical subject matters are a rich source of felicitous conspiracy conditionals:

- (27) If you're in the good case, then you're not a handless BIV.
(28) If you don't know you have hands, then you're in the bad case.
(29) If nature is uniform, then induction is a reliable means of forming true beliefs.
(30) If the next time I drop this pen it floats in the air instead of falling, then either it's a very special pen or nature may not be so uniform after all.

But again, if we know anything by the moderate invariantist's lights, then we know that we know we have hands and that pens fall when they are dropped.²⁴

Here is the upshot. We have reason to believe IGNORANCE is true and that each of the above conspiracy conditionals is felicitously and canonically assertable in at least some contexts. Since moderate invariantism tells us that if the proposition expressed by $\ulcorner S$ knows whether $p \urcorner$ is false in one context it is false in every context, and since IGNORANCE tells us that if S 's canonical assertion of \ulcorner If p , $q \urcorner$ is felicitous in a context then the proposition expressed by $\ulcorner S$ knows whether $p \urcorner$ is false in that context, IGNORANCE plus the facts

²⁴One philosophical case I find particularly sharp draws on some recent papers (cited below) on epistemic puzzles concerning unobserved tosses of fair coins. Suppose you know a fair coin is about to be flipped 1,000 times. Supposing the distribution of heads/tails *in fact* ends up being relatively normal, do you know prior to the coin's being flipped that it won't land heads all 1,000 times? There are persuasive arguments suggesting that the moderate invariantist must say 'Yes'—see, e.g., (Bacon, 2014, §1) and (Dorr *et al.*, 2014). Supposing these arguments are sound, it follows that these indicatives are conspiracy conditionals:

- (31) a. If this fair coin lands heads 1,000 times in a row, then most onlookers will be quite confident that the coin is double-headed.
b. If this fair coin doesn't land heads 1,000 times in a row, then anyone who bet that it would will lose some money.

about the conspiracy conditionals look to eliminate large swaths of what the moderate invariantist takes us to know. That is the paper's core argument.

I see three ways the moderate invariantist could respond. First, accept IGNORANCE as stated and argue that the felicitous uses of conspiracy conditionals are invariably non-canonical. Second, reject IGNORANCE in favor of a non-knowledge-centric theory of the pragmatics of indicative conditionals. Or third, again accept IGNORANCE as stated, but give an error theory of our inclination to judge felicitous canonical uses of conspiracy conditionals. Each of these replies will be critically assessed in §§5, 6, and 7 respectively.

5 Must conspiracy conditionals be non-canonical?

Supposing the felicitous uses of conspiracy conditionals are *always* non-canonical, then IGNORANCE in combination with our judgments about those uses is no threat to the moderate invariantist's theory of knowledge. How might one defend this supposition? First, one could try to assimilate the relevant uses of conspiracy conditionals with the *canonically* non-canonical uses of the indicative conditional. That is to say, one could try and classify the natural uses of §4's conspiracy conditionals as kinds of echoing, Dutchman, or concessive uses. Alternatively, one could argue that conspiracy conditionals give rise to a *sui generis* kind of non-canonical use of the indicative conditional—a yet further exception to the general prohibition on asserting conditionals whose antecedents are known. I will argue that neither option is plausible.

5.1 Conspiracy conditionals as canonically non-canonical

Beginning with the first option, should we think that the natural uses of conspiracy conditionals are always echoing, Dutchman, or concessive uses?

The short answer is: No, we shouldn't. With only minimal amounts of reflection one should be able to see that ordinary uses of (e.g.)

(1) If Oswald didn't shoot Kennedy, then someone else did.

(26) If Federer didn't defeat Murray in straight sets, then this report is inaccurate.

are not invariably echoing, Dutchman, or concessive uses. This is not to say they *can't* be. No one should deny that the same conditional can be used in all sorts of different ways. It's just that there are plenty of normal contexts in which, intuitively, one's intention in asserting a conditional like (1) or (26) is to tell one's interlocutors something informative about the subject of conversation, rather than to present evidence for what one knows or accommodate their distrustfulness.

A case or two will help sharpen the point. Start with our paradigm conspiracy conditional:

(1) If Oswald didn't shoot Kennedy, then someone else did.

It goes without saying that (1) is not ordinarily used as a Dutchman. Must it be used in an echoic or concessive manner? Well, given that Oswald shot Kennedy (and that we context-invariantly know as much according to moderate invariantism), if it is to be used in an echoing way it has to be in something like a modus tollens arguments for the negation of its antecedent—i.e., for the conclusion that Oswald *did* shoot Kennedy. But such an argument looks baffling:

If Oswald didn't shoot Kennedy, then someone else did. We know that someone else didn't shoot Kennedy. ?? Therefore, Oswald shot Kennedy.

It is hard to imagine circumstances in which this reasoning accomplishes anything. Again, this is not to say that there are *no* possible circumstances in which it does—surely some could be conjured up. The point is just that the typical circumstances in which one would be inclined to assert or otherwise endorse (1)—like those first encounters with the example early in one's philosophical career—one will almost surely not be doing so in this way. It is thus not plausible that the default uses of (1) are echoing uses.

Concessive uses of (1) are more palatable. This is because a concessive use of (1) will be appropriate whenever one's interlocutors will not accept that you know Oswald shot Kennedy, but will accept that you know that *someone* did—circumstances which are not particularly difficult to imagine. But what *is* difficult to imagine is that these are the *only* circumstances in which an assertion of (1) is appropriate. In deciding to assert (1), one needn't think something like 'Well of course *I* know that Oswald shot Kennedy, but since these people won't believe me I'll have to settle for "If Oswald didn't shoot Kennedy, then someone else did"'. No—*pace* the error theories to be discussed in §7—all goes perfectly well if one thinks 'I'm not sure whether Oswald shot Kennedy; but what I am sure of is that if Oswald didn't, someone else did' or 'Hmm—if Oswald didn't shoot Kennedy, then someone else did. I wonder who else might have shot him...' or what have you. The corresponding uses look to be straightforwardly canonical.

The same points go through with the other examples. Recall the scenario where you read in the perfectly reliable local newspaper that Federer defeated Murray in straight sets, as well as its conspiracy conditionals:

(25) If Federer defeated Murray in straight sets, then this report is accurate.

(26) If Federer didn't defeat Murray in straight sets, then this report isn't accurate.

There are plenty of ordinary contexts in which either may be asserted felicitously. Are such uses invariably non-canonical?

Again the answer is No. The natural uses of (25) and (26) are obviously not Dutchmen. They are equally obviously not concessive: prefixing an 'even' to the beginning of either

conditional makes it infelicitous. So if their natural uses are (canonically) non-canonical, it is because their natural uses are echoing uses. But it is highly implausible that an assertion of either (25) or (26) always or even *typically* has an echoing function. Try to imagine non-peculiar circumstances in which your total evidence regarding what happened in the match is what you read in the report, and in which you would assert (25) as part of an argument for the report's accuracy, or (26) as part of an argument for Federer's having defeated Murray in straight sets. If you know Federer defeated Murray in straight sets it's because you read the report saying as much. So even if you know that Federer's defeating Murray in straights entails (given the rest of what you know) that the report is accurate, the evidential relation between those propositions makes (25) an unconvincing way of echoing your knowledge, to say the least. Likewise, although arguing from the fact that a reliable outlet reports that Federer defeated Murray in straights to the conclusion that Federer in fact defeated Murray in straights is sometimes a reasonable thing to do, what is not a reasonable thing to do is to argue for it in the roundabout, modus tollens style way that would make (26) an appropriate thing to assert. But all this is just a fancy way of saying something that should be obvious: neither (25) nor (26) is the kind of thing you'd expect a speaker to assert if she were trying to provide or rehearse evidence that Federer defeated Murray in straight sets. (25) and (26) may have echoing uses in some contexts, but it strains credulity to think they have them in *every* context. But that is the claim the moderate invariantist needs to make if she wants her theory to be able to accommodate IGNORANCE.

I will spare the reader similar analyses of the other examples discussed in §4. The analysis of the Oswald and Federer conspiracy conditionals should be enough to establish high confidence in this subsection's modest thesis: that the felicitous uses of conspiracy conditionals need not be (and often are not) among the familiar non-canonical uses of the indicative conditional. If they are deserving of some special non-canonical status, it must be of a *sui generis* kind.

5.2 Conspiracy conditionals as non-canonically non-canonical

A proponent of this style of response has it that §3's taxonomy of non-canonical uses of the indicative conditional is incomplete: in addition to **echoing**, **Dutchman**, and **concessive** uses, there are also **conspiratorial** uses, or whatever you want to call them.

What's the argument in favor of expanding the domain of non-canonical uses of the indicative conditional to cover the natural uses of conspiracy conditionals? Well, presumably, it's that we want both IGNORANCE and moderate invariantism to be true, and we can't have that unless the natural uses of conspiracy conditionals form their own category of non-canonical use.

Aside from its dialectical awkwardness, I think there are strong reasons to be suspicious of this line of thinking. A unifying thread among the familiar kinds of non-canonical

uses of the indicative conditional is that each has a *pragmatic* explanation for its existence. Echoing uses arise because sometimes it's important that speakers highlight inferential connections; Dutchman uses are just echoing uses with sarcasm, and so presumably arise for whatever reason sarcastic uses of language arise more generally; and concessive uses arise because sometimes speakers have to deal with interlocutors that are insufficiently trusting.

But the plausible explanations of the existence of conspiratorial uses cannot be like this. And that's because conspiracy conditionals may be used felicitously even in situations in which all the standard pragmatic pressures are present—pressures to assert the strongest thing you are in a position to assert, keep it brief, etc. As far as I can tell, then, the story that explains what makes the natural uses of conspiracy conditionals non-canonical will have to be an *epistemic* one. That is, it will have to be a story that says that despite the fact that there is nothing pragmatic that would encourage a speaker to assert 'If p , q ' rather than p , there is something about the speaker's epistemic state that makes the assertion of 'If p , q ' preferable to an assertion of p .

The problem with such a story is that it is unclear what distinguishes (i) a view that says there are uses of the indicative conditional that are non-canonical because they arise in contexts with non-standard *epistemic* (rather than pragmatic) constraints on assertion from (ii) a view that says that there are uses of the indicative conditional that are counterexamples to IGNORANCE. One who accepts the former but denies the latter risks trivializing IGNORANCE by turning it into a principle that says that conditionals with known antecedents are assertable except when they're not.

If the moderate invariantist wants to stay in the business of taking our judgments about conspiracy conditionals at face value, then rather than invent *ad hoc* categories of non-canonical uses of the indicative conditional, she should simply take these conditionals to show that IGNORANCE must be false. If it is non-negotiable that the proposition expressed by 'S knows whether p ' is true in every context, and if nonetheless there are contexts in which S may felicitously assert 'If p , q ' in exactly the same manner she may assert any other ordinary indicative conditional, then—contra IGNORANCE—there must be no *general* prohibition on asserting 'If p , q ' in the canonical way when you know whether p .

6 Replacing IGNORANCE?

So: assuming error theories of our judgments about the felicitous assertion of conspiracy conditionals are off the table, the moderate invariantist must reject IGNORANCE.

Here are two ways she might go about doing that. First, she could reject IGNORANCE while offering no suggestion as to how we might try to account for the empirical and theoretical considerations that compelled us to posit it in the first place. In other words, she

could take conspiracy conditionals (and perhaps also the non-canonical uses of the indicative conditional) to be cause to abandon the prospects of *any* general theory of the ignorance implicatures of the indicative conditional. Alternatively, she could take IGNORANCE to be tracking *something*, but deny that it's knowledge that determines the felicity conditions of the canonical use of an indicative conditional. Instead, she will try to find some other epistemic relation—perhaps invariant, perhaps context-sensitive—and argue that the facts about that relation (in context) determine the felicity conditions.

I am not going to consider the first way of rejecting IGNORANCE. Even if IGNORANCE is ultimately false, the move to the conclusion that there are no interesting generalizations in its vicinity is at best a measure of last resort. I will thus focus on the second way of rejecting IGNORANCE, which is to replace it with a comparable principle stated in terms other than knowledge. Consequently, I will take it for granted that everyone agrees that some version of the following placeholder principle is true:

OPENNESS For canonical uses of the indicative 'If p , q ': In every context c : if 'It isn't open for S whether p ' is true in c , then ' S may *not* felicitously assert the indicative conditional 'If p , q ' is true in c .

OPENNESS tells us that whether S may felicitously assert 'If p , q ' in the canonical way depends on whether S stands in a placeholder epistemic relation to p . In particular, if p isn't *open* for S , then S cannot assert 'If p , q '. This section will look at three theories of the *being open* relation, and assess whether they let OPENNESS give a satisfactory explanation of the considerations §2 raised in favor of IGNORANCE (*inter alia*). They are: (i) theories that take the underlying epistemic relation to be invariantly something like *belief*, (ii) theories that take the underlying epistemic relation to be invariantly something like *certainty*, and (iii) theories that give 'is open' a contextualist semantics distinct from the semantics of 'knows'. In each case it will be argued that a knowledge-theoretic conception of OPENNESS is preferable to the alternative.

6.1 Against belief-theoretic conceptions of openness

On a belief-theoretic conception of openness, a proposition p is open to S just in case S neither believes p nor $\neg p$. Understood in this way, OPENNESS says that whenever one believes that p , one should assert neither 'If p , q ' nor 'If $\neg p$, q '.²⁵

There are two fairly obvious problems for the belief-theoretic conception of OPENNESS, the first particular to the moderate invariantist, the second more general. First, on most standard views of knowledge, S 's knowing that p entails S 's believing that p . So if contexts in which S 's assertion of the conspiracy conditional 'If p , q ' is felicitous are contexts in

²⁵ It is to be implicit through this section that unless stated otherwise we are restricting attention to canonical uses of the indicative conditional.

which S fails to believe either that p or that $\neg p$, then on any of these accounts of knowledge S will automatically fail to know whether p too. The problem for the moderate invariantist thus persists.

Second, the view that p is open to S iff S neither believes that p nor that $\neg p$ is empirically inadequate. Sentences of the form ‘I believe that p , but if $\neg p$, q ’ sound perfectly natural, while sentences of the form ‘I know that p , but if $\neg p$, q ’ and ‘ p , but if $\neg p$, q ’ sound terrible:

- (32) a. I believe that Jim is at home listening to *The Mollusk*, but if he’s at the office then he probably isn’t listening to music at all.
 b. ?? I know that Jim is at home listening to *The Mollusk*, but if he’s at the office then he probably isn’t listening to music at all.
 c. ?? Jim is at home listening to *The Mollusk*, but if he’s at the office then he probably isn’t listening to music at all.
- (33) a. I believe that white is winning in this position, but if it’s not then the reason why is beyond my knowledge of chess.
 b. ?? I know that white is winning in this position, but if it’s not then the reason why is beyond my knowledge of chess.
 c. ?? White is winning in this position, but if it’s not then the reason why is beyond my knowledge of chess.

Given a belief-theoretic implementation of OPENNESS, (32a) nor (33a) should both have the feel of their (b-c) counterparts, when in fact they do not.

6.2 Against certainty-theoretic conceptions of OPENNESS

The second problem with the belief-theoretic conception of openness suggest that epistemic relations that are generally *weaker* than knowledge will fail to give the moderate invariantist the OPENNESS she needs. This should come as no surprise: the phenomenon made salient by conspiracy conditionals is that there are some contexts in which one has to do *more* than know a proposition or its negation to get the infelicity of the relevant conditional. So a view that ties openness to an epistemic status that is even easier to achieve than knowledge—which is exactly what the belief view does—is not likely to be of much help.

In light of all this, a natural thought is to use an epistemic relation that is *stronger* than knowledge to determine which possibilities are open for a subject. Let us use *certainty* as a placeholder for such a relation. One candidate gloss has it that S is certain that p just in

case S has credence 1 that p .²⁶ Another possible gloss has it that S is certain that p just in case S stands to p in whatever relation is supplied by the skeptical invariantist's preferred theory of knowledge. That is to say, S is certain that p just in case S knows that p by the lights of the skeptic.²⁷

We can leave these and other glosses on 'is certain' more or less unanalyzed. What matters is just that the kinds of propositions to which we stand in such relations are generally quite limited—i.e., that an ordinary person S is *not* certain of propositions like:

(P₁) That S has hands.

(P₂) That Donald Trump is the president of the United States.

(P₃) That Oswald shot Kennedy.

So long as the conception of certainty can deliver this, there is no problem reconciling OPENNESS with our intuitive judgments about conspiracy conditionals.

Still, the view faces other significant challenges. The main one is to explain the compatibility of the following three bits of data. First, that there are some contexts in which ordinary speakers may felicitously assert sentences that express any of (P₁)–(P₃). Second, that there are also some contexts in which speakers may felicitously assert the corresponding conspiracy conditionals. And third, that there are no ordinary contexts in which speakers may felicitously assert the *conjunction* of the two—i.e., sentences of the form 'p but if ¬p, q' or 'p and if p, q':

- (34) I'll try to hit a kick serve out wide next point. ?? If I have hands, then I'll make sure to put extra top spin on it.
- (35) Tomorrow we're going to the White House to see the president. ?? If there is a president of the United States, it's Trump.
- (36) Well, everyone knows Oswald shot Kennedy. ?? But if he didn't shoot Kennedy then someone else did.

As far as I can tell, the certainty-theoretic conception of openness has no good explanation of the conjunction of these facts. In short: her conception of openness forces her to give a non-unified treatment of the norms governing the felicitous use of conditionals and the felicitous use of bare declaratives, but the proper explanation of the context-invariant badness of speeches of the form 'p but if ¬p, q' requires that those norms be unified.

²⁶ Perhaps credence 1 is too strong. If one has the view that there is nothing a rational person has credence 1 in, then the credence 1 conception of OPENNESS has it that no indicative conditional is ever ruled out—not even canonical uses of conditionals like 'If I do not exist, then...'. Those who are nervous about such things should interpret credence 1 as credence approximately equal to 1; it won't make a difference to the arguments of this section.

²⁷ See, e.g., Unger (1975) and Hawthorne (2004, §§3.1-3.5) for discussion of what the skeptic's theory of knowledge amounts to.

The longer version of the argument starts with two obvious points. First, the certainty conception of OPENNESS cannot by itself explain the badness of any of (34)–(36). No ordinary person is certain (in our technical sense) that they have hands or that Oswald shot Kennedy. These propositions and their negations will thus be open in any context, and so acceptable by the lights of OPENNESS. Of course, this is not yet a strike against the certainty conception of OPENNESS, for the ambition of the principle has only ever been to provide *necessary* conditions on the felicitous assertion of an indicative conditional. But it is important to have this observation in the background.

The second obvious point is that those who understand the *being open* relation in certainty-theoretic terms cannot endorse:

ASSERTION In every context c : if $\ulcorner S$ may (epistemically) assert $p \urcorner$ is true in c , then \ulcorner It isn't open for S whether $p \urcorner$ is true in c .

What ASSERTION says, in words, is that a *necessary* condition on S 's asserting p is that p not be open for S . The obvious reason this principle is unacceptable for those who understand openness in terms of the absence of certainty is that if it were true, it would follow that no one has ever felicitously asserted anything, barring assertions of tautologies and the like. But we're taking it as a datum that sentences expressing mundane propositions like (P₁)-(P₃) have been felicitously asserted. So ASSERTION is a non-starter when understood certainty-theoretically.

With these two points in mind we can now state the problem. Were both OPENNESS and ASSERTION true, then there would be an easy explanation of the badness of (34)–(36): in any context in which $\ulcorner S$ may assert $p \urcorner$ is true, p is not open to S ; but if p is not open to S , then $\ulcorner S$ may assert \ulcorner If $\neg p$, $q \urcorner$ is false. So there will be no context in which S may assert either $\ulcorner p$ and if p , $q \urcorner$ or $\ulcorner p$ but if $\neg p$, $q \urcorner$. That is to say, OPENNESS plus ASSERTION entails COORDINATION:

COORDINATION In every context c : if $\ulcorner S$ may (epistemically) assert $p \urcorner$ is true in c , then $\ulcorner S$ may *not* felicitously assert either \ulcorner If p , $q \urcorner$ or \ulcorner If $\neg p$, $q \urcorner$ is true in c .

And COORDINATION entails that (34)–(36) are always unassertable. So those who accept both OPENNESS and ASSERTION can explain COORDINATION, which in turn lets them explain our judgments about assertions of the form $\ulcorner p$ and if $(\neg)p$, $q \urcorner$.

But those who use a certainty-theoretic conception of openness cannot avail themselves of such an explanation, for by their lights ASSERTION is false. Instead, they must make COORDINATION true by *fiat* to get the results they want. Why by fiat? Because on this kind of view, the pragmatics of indicative conditionals tracks one kind of epistemic relation (either credence 1 or the skeptic's knowledge), the pragmatics of assertion another (the moderate invariantist's knowledge? belief?), and yet still, for some unexplained reason,

there is no context in which one is permitted to assert both p and $\lceil \text{If } (\neg)p, q \rceil$. Such a view looks worryingly ad hoc.

6.3 Contextualism about openness

The certainty-theoretic conceptions of OPENNESS cannot predict COORDINATION. The lesson to take from this is that if one wants take seriously our ordinary judgments about the assertion of conspiracy conditionals and their bare antecedents, then one is going to have to be a *contextualist* about whatever relation gets plugged into OPENNESS.²⁸ If the relation is invariant and stricter than knowledge, then although we can explain why conspiracy conditionals sometimes sound good, we cannot explain why they sometimes sound bad. And if the relation is invariant but no stricter than knowledge, then although we can explain why conspiracy conditionals sometimes sound bad, we cannot explain why they sometimes sound good. So the relation needs to vary from context to context. Thus, if one thinks that the best versions of OPENNESS and ASSERTION are knowledge-theoretic—i.e., one likes IGNORANCE and KNA:

IGNORANCE For canonical uses of the indicative $\lceil \text{If } p, q \rceil$: In every context c : if $\lceil S$ knows whether $p \rceil$ is true in c , then $\lceil S$ may *not* felicitously assert the indicative conditional $\lceil \text{If } p, q \rceil$ is true in c .

KNA In every context c : If $\lceil S$ may (epistemically) assert $p \rceil$ is true in c , then $\lceil S$ knows that $p \rceil$ is true in c .

—then one is going to have to be a contextualist about knowledge, and *ipso facto* not a moderate (or skeptical) invariantist. Otherwise one is going to have to find some other context-sensitive notion to play the *being open* role.

So: supposing contextualism about knowledge is off the table—as it obviously is for the moderate invariantist—what are the prospects of finding a contextualist, non-knowledge-centric notion capable of delivering both OPENNESS and ASSERTION?

6.3.1 Openness in terms of common ground

One tempting line here is to connect openness to the notion of the *common ground*: the set of possibilities consistent with all that has been mutually presupposed and asserted by a context's conversational participants.²⁹ On this view, a proposition p is open in a context c

²⁸ This may be too quick. Nothing that has been said so far rules out that a *sensitive invariantist* conception of knowledge (e.g., that of Hawthorne (2004) or Stanley (2005)) could combine with IGNORANCE to deliver COORDINATION and our intuitive judgments about conspiracy conditionals. The argument against that package of views will come in §8.

²⁹ The vague characterization just given suffices for present purposes, but see, e.g., Stalnaker (1999, 2002) and the citations therein for discussion.

iff neither p nor $\neg p$ is entailed by the conversational common ground associated with c . So given OPENNESS, the assertion of an indicative conditional $\lceil \text{If } p, q \rceil$ is felicitous in c only if the answer to the question of whether p is not common ground in c .³⁰

Despite whatever intuitive appeal it has, the common ground-theoretic account of OPENNESS faces some severe challenges. For starters, it fails to explain some of even the most basic data motivating generalizations like OPENNESS. Here I have in mind examples like (2) (? ‘If I had a bagel for breakfast this morning, then I went out for brunch with Jim’). As was argued in §2, an assertion of (2) will in most ordinary contexts be puzzling. But there should be nothing strange about it by the lights of OPENNESS (on the assumption that ‘being open’ is a matter of being compatible with the conversational common ground). After all: neither the proposition that Jane had a bagel for breakfast nor that proposition’s negation is common ground in the imagined scenario.

Perhaps more worrying, the view fails to explain the badness of speeches of the form $\lceil \text{I know whether } p. \text{ If } (\neg)p, q \rceil$. Consider, e.g.:

- (37) I know whether Oswald shot Kennedy. ?? If Oswald shot Kennedy, then he probably acted alone.
- (38) I know whether Trump is the president. ?? If Trump isn’t the president then Pence is.

Given OPENNESS (again understood common ground-theoretically), (37) and (38) should be felicitous in any context in which $\lceil \text{It is open for } S \text{ whether } p \rceil$ is true—i.e., contexts in which the speaker’s knowing whether p (and having asserted as much) does not suffice for its being common ground whether p . Thus, since there seem to be no contexts in which either (37) or (38) is felicitous, it better be that an assertion of $\lceil S \text{ knows whether } p \rceil$ has the effect of making common ground whichever of p and $\neg p$ is true.

But how is *that* supposed to happen? Unlike $\lceil S \text{ knows that } p \rceil$, $\lceil S \text{ knows whether } p \rceil$ neither entails nor presupposes that p . So even if the facts about whether p is common ground tell us about the assertability of p and $\lceil \text{If } \neg p, q \rceil$, they tell us nothing about the assertability of $\lceil S \text{ knows whether } p \rceil$. This means, among other things, that the proponent of the current view cannot claim that because $\lceil \text{It is open to } S \text{ whether } p \rceil$ is true in context, $\lceil S \text{ knows whether } p \rceil$ must be either false or unassertable in context. Likewise, one cannot claim that because $\lceil \text{It isn’t open to } S \text{ whether } p \rceil$ is true in context, $\lceil S \text{ knows whether } p \rceil$ must be true in context too. Thus, to explain the badness of speeches like (37)–(38), one

³⁰Note, however, that ASSERTION so understood is obviously a non-starter: if one can only assert that p if it is common ground that p , then assertion cannot serve its primary (perhaps even constitutive) purpose of eliminating possibilities from the common ground. And this means a proponent of the common ground view cannot get COORDINATION. But this isn’t obviously so bad, for she can get something pretty close to it no matter her theory of ASSERTION. This is because the context-invariant badness of sentences like $\lceil p \text{ but if } \neg p, q \rceil$ follows from the fact that an assertion of p has the effect of making it common ground that p , which renders $\lceil \text{If } \neg p, q \rceil$ infelicitous.

who accepts the common ground-theoretic conception of OPENNESS must posit further, ad hoc connections between ‘is open’ and ‘knows’. Something like: whenever ‘It is open to S whether p ’ is true in context, ‘ S knows whether p ’ is false (or at least unassertable) in context—even if neither p nor $\neg p$ is common ground. But I see no coherent methodological perspective from which one could sensibly stipulate this.

6.3.2 Epistemic contextualism about openness

The lesson to take from the foregoing is that the truth of ‘It is open to S whether p ’ in c is going to have constitutive ties to the *epistemic* facts about S (in c), and not merely to facts about what S (and her conversational partners) have decided to presuppose or assert in c . And on that front we’ve already considered three epistemic relations—invariantism about knowledge, belief, and certainty—and found them wanting. Thus, to avoid the argument for contextualism about knowledge, the moderate invariantist will need to find some other appropriately context-sensitive epistemic notion to play the openness role. This forces her to choose between two options, broadly construed: either (i) connect the judgments that motivate generalizations like IGNORANCE and KNA to *familiar* context-sensitive epistemic notions other than ‘knows’; or (ii) invent some new epistemic vocabulary to model the workings of ‘is open’.

I am skeptical that there are familiar epistemic notions (other than ‘knows’) that could plausibly characterize what it is for a proposition to be open in the sense relevant to OPENNESS and ASSERTION. I am also skeptical that one who invents epistemic vocabulary can do much better than to treat ‘is open’ as an unanalyzed primitive. But regardless of whether I’m right on those matters (and how bad it is if I am), I want to raise two worries—one empirical, one conceptual—that should apply to any view in the vicinity here.

The empirical worry is that by severing the connection between ‘knows’ and ‘is open’ (and *ipso facto* denying principles like IGNORANCE and KNA), it will be difficult to explain quite a lot of the intuitive data that compelled us to posit those principles in the first place. If there is not some general principle to the effect of ‘Assert p only if you know that p ’, then what explains the naturalness of certain stock challenges to assertions, e.g.: ‘How do you know that?’, ‘You don’t know that’, etc.? Likewise: as was observed in §2, when someone asserts ‘If p , q ’ in a situation in which we’d expect that person to know whether p , there is an inclination to respond with puzzlement along the lines of ‘What do you mean if p ? Don’t you know whether p ?’. Is this inclination the result of a mistake? Or does it just so happen that by coincidence, whatever epistemic relation is denoted by ‘is open’ in standard contexts is the same as (or suitably related to) the one denoted by ‘knows’? I see no straightforward answers to questions like these.

The conceptual worry is more of a *tu quoque* against the moderate invariantist; but given that she more than anyone has reason to want to deny principles like IGNORANCE (or so

I've argued), having it on the table will help clarify some of IGNORANCE's attraction. The worry is this. One who gives a theory of the 'is open' of OPENNESS and ASSERTION in terms of something other than 'knows' must reject a hypothesis many moderate invariantists are likely to find antecedently appealing. It is the hypothesis that *knowledge* is what determines (e.g.) the norms of assertion and the ignorance implicatures of indicative conditionals. In rejecting this hypothesis, the moderate invariantist robs the concept of knowledge of much of its explanatory attractiveness. Worse, in positing this basic yet (apparently) deeply engrained epistemic notion, the moderate invariantist faces a renewed threat from all sorts of other knowledge-obviating views. Given that *not being open*—whatever that amounts to—is what governs assertion and the pragmatics of indicative conditionals, one might wonder whether it could have a similar role in explaining our judgments about (e.g.) evidence and action. Better to tie OPENNESS and ASSERTION to a contextualist account of 'knows', and thereby allow knowledge to maintain much of its explanatory power.

7 An error theory about conspiracy conditionals?

We know from §4 that given IGNORANCE, felicitous canonical uses of conspiracy conditionals pose a direct challenge to moderate invariantism about knowledge. And we know from §5 that the natural uses of conspiracy conditionals are sometimes (indeed often) canonical, and from §6 that if there's a true generalization about the ignorance implicatures of indicative conditionals, it's a knowledge-theoretic one like IGNORANCE. Supposing the arguments of these sections are sound, then conspiracy conditionals really do present a strong challenge to moderate invariantism about knowledge. Indeed, for the view to survive, it must be that our inclination find canonical uses of conspiracy conditionals felicitous is the result of a systematic error. This section will investigate the plausibility of an error theory about conspiracy conditionals.

The task of explaining away ordinary judgments about knowledge is familiar territory for the moderate invariantist. It is well known that the facts about our ordinary judgments about knowledge *ascriptions*—sentences of the form 'S knows whether p'—are at least somewhat in tension with the moderate invariantist's semantics. Accounting for these facts has forced her to develop error theories. One might wonder whether those same error theories could (and if so, should) be called upon to explain our intuitions about conspiracy conditionals. This section will give two arguments suggesting that our judgments about conspiracy conditionals ought not be accounted for in this way—one quick, one long.

The quick argument is this. Even if the plausibility of the moderate invariantist's preferred error theory isn't fundamentally affected by the shift from knowledge ascriptions to indicative conditionals, it is still *yet another* bit of language that competent speakers are alleged to regularly screw up. The wider the net of the error theory, the more seriously we

should consider the possibility that the semantics is failing to model the right thing. This is especially true when the practices that fall into the domain of the error theory look as innocuous as they do here. Take the following conspiracy conditionals, for example:

- (1) If Oswald didn't shoot Kennedy, then someone else did.
- (25) If Federer defeated Murray in straight sets, then this report is accurate.
- (29) If nature is uniform, then induction is a reliable means of forming true beliefs.

If IGNORANCE and moderate invariantism are both true, then there is *no* context in which any of these is an appropriate thing to say. But our inclination to think and act otherwise just doesn't look like the kind of practice that calls for therapy.

That's the quick argument. The longer argument is that our practice of asserting conspiracy conditionals looks to be particularly difficult to explain away error-theoretically. That is to say, the kinds of factors cited by the moderate invariantist in her explanation of the variability of our intuitions about knowledge ascriptions do not extend naturally to many of the situations in which conspiracy conditionals seem assertable. Moreover, it is unclear whether there are any plausible amendments that could be made to such a story to account for the new recalcitrant data. But arguing for this will first require a brief overview of the state of the dialectic surrounding the old error theories. So with that...

7.1 Interlude: the old dialectic

Let's say that the *purely epistemic* factors that go into determining whether a subject *S* knows a proposition *p* are exhausted by *S*'s beliefs and evidence—where 'evidence' is to be understood in an intuitive, theory-neutral way—as well as the truth of the matter about *p*. As is well known, there can be pairs of situations that agree on the purely epistemic factors relevant to the relationship between a subject *S* and proposition *p*, and yet in which the truth of the *knowledge ascription* 'S knows that *p*' is nonetheless judged to differ.³¹ For example, it is not at all difficult to find contexts in which for any of (P₁)-(P₃), intuitively 'S knows that (P_{*n*})' is true. However, it is also not at all difficult to find contexts in which *holding the purely epistemic facts about S fixed*, it is the negation of 'S knows that (P_{*n*})' that is intuitively true. Contexts in which the following are natural tend to do the trick:

- (39) Unless *S* has discovered some way of proving that she isn't a handless BIV, *S* doesn't know that she has hands.
- (40) *S* doesn't know that Trump hasn't had a fatal heart attack in the past few moments, so *S* doesn't know that Trump is still the president of the United States.

³¹ For just a small sampling of the massive literature on this phenomenon, with a particular focus on work that offers ordinary language evidence in favor of variants to moderate and skeptical invariantism, see, e.g., Austin (1946); Cohen (1986, 1999); DeRose (1992, 1995, 2002); Hawthorne (2004); Lewis (1996); Schaffer (2005b, 2007); Schaffer & Szabo (2013); Stanley (2005).

- (41) There are too many unanswered questions surrounding Oswald's involvement in Kennedy's assassination for *S* to know for sure that Oswald shot Kennedy.

Thus, marginal cases aside, the following empirical fact is common ground to those who theorize about knowledge: that there can be pairs of situations in which all the purely epistemic factors relevant to *S* and *p* are fixed and in which the truth of the knowledge ascription 'S knows that *p*' is judged to vary by competent speakers of the language. By extension, it is common ground that's one theory of knowledge must account for this fact. This is not to say that it must codify it into the semantics for 'knows' itself; it's just that if the theory takes competent speakers to be judging or speaking falsely, it better explain why this happens.

Enter the moderate invariantist's error theory. The first half of the error theory is the empirical half. It identifies some psychological mechanism—which we'll call *X*—and observes that the presence of *X* correlates systematically with a tendency to deny the ascriptions that, by the lights of the theory, ought in every context be affirmed. The second half is the exculpatory half. It provides reasons to believe that the presence of *X* explains these variations in a way that relieves the pressure to have one's semantics for 'knows' account for them.

What is *X*, the mechanism the error theory finds predictive of the variations it seeks to explain away? There are, in general, two kinds of sub-mechanisms discussed in the literature: *stakes* and *salience*.³² Both can be further broken down into cases affecting the ascriber (*A*) and cases affecting the subject (*S*).³³ This gives us four factors in total: what the stakes are for *A*, what the stakes are for *S*, what's salient to *A*, and what's salient to *S*. We'll say more about what these factors amount to in a moment. For present purposes it is enough if one's understanding of the nature of these mechanisms is mostly exhausted by one's knowledge of the ordinary meaning of the words 'stakes' and 'salience'.

The *empirical* half of the error theory says that as you vary the stakes or what's salient (perhaps for *A*, perhaps for *S*), you should expect *A*'s judgments about the truth of 'S knows that *p*' to vary. In particular, as what is at stake on whether *p* becomes higher, and as (metaphysical) possibilities in which $\neg p$ become more salient (again perhaps to *A*, perhaps to *S*), *A* will become more inclined to judge that 'S knows that *p*' is false.³⁴

³² The literature here is truly massive. See, e.g., Adler (2006); Bach (2005); Buckwalter & Schaffer (2015); Fantl & McGrath (2002, 2009); Hawthorne (2004, §4.2); May *et al.* (2010); McGrath (2007); Nagel (2008, 2010); Reed (2010); Roeber (2018); Rysiew (2001); Schaffer (2005a); Stanley (2005); Weatherson (2012); Williamson (2005). Note that many of these authors—in particular those who defend subject-sensitive accounts of knowledge—build certain aspects of *X* into the semantics of 'knows' itself. As we will see in §8, whether *X* goes into one's semantics or one's error theory makes no difference to the problem posed by conspiracy conditionals.

³³ In the case of first-personal knowledge ascriptions $A = S$.

³⁴ The exact details of the weights each of these four factors ought to be given is an open question. As we will see shortly, however, whatever the answer will ultimately make no difference to the dialectic of this paper. What matters is that *X* is exhausted by the *disjunction* of these four factors. And to my knowledge that much has been taken to be uncontroversial in the literature.

The exculpatory half of the error theory then argues that we have reason to expect that these sorts of variations in stakes and salience should cause *errors* in ordinary attributions of knowledge. For two prototypical examples of this half of the error theory (as well hints about the nature of stakes and salience), see [Williamson's \(2005, p. 226\)](#) line:

[Possibilities of error] may be psychologically salient because the practical costs of error are high for the subject or the ascriber, or simply because they have been evoked in vivid and convincing detail. One effect of fictional violence on television is to make viewers overestimate their chances of being the victims of violent crime: they suffer an illusion of danger. Might not an illusion of epistemic danger result from exposure to lurid stories about brains in vats, evil demons, painted mules, or gamblers who bet the farm?

As well as [Hawthorne's \(2004, p. 164\)](#):

[Psychologists] emphasize the role played by the 'availability' heuristic as a distorting influence on our judgments of risk: in many cases, our estimation of the likelihood of an event is affected by the ease with which we can recall or imagine it...Applied to the issue at hand, the availability heuristic may help to explain our tendency to skeptical overprojection. When certain non-knowledge-destroy counterpossibilities are made salient, we overestimate their real danger; as a result, we may find ourselves inclined to deny knowledge to others in cases where there is in fact no real danger of error.³⁵

So, to summarize: the moderate invariantist's error theory of our intuitions about knowledge ascriptions is composed of two hypotheses. First, that as a matter of empirical fact, mechanism *X*—i.e., variations in what is at stake for or salient to ascriber and/or subject—is what explains the variability in our intuitive judgments about the truth of various mundane knowledge ascriptions. And second, that the presence of *X* is the kind of thing that can sympathetically explain why otherwise competent speakers so often err in their use of 'knows'.

7.2 The new dialectic

Let us suppose that the error theory just described is a contender as concerns our intuitions about knowledge ascriptions. Should we think that it works as an error theory of our judgments about *conspiracy conditionals*? Well, that depends on whether the two components of

³⁵ [Nagel \(2010\)](#) offers an extended criticism of Hawthorne and Williamson's use of the availability heuristic in their error theories. Her preferred account appeals to the phenomenon of 'epistemic egocentrism', whereby people fail to suppress privileged information (i.e. what they know or are concerned with) in evaluating the judgments of others. The complaints about Williamson and Hawthorne's error theories will apply just as much to hers.

the error theory about the former can be transposed into an error theory about the latter. And on that issue I think that the empirical half of the error theory is almost surely inadequate, and that to the extent that it can be recovered, it is only in virtue of compromising the exculpatory half.

The empirical half of the moderate invariantist's error theory associates variations in our judgments about ascriptions with changes in stakes and salience. Since by IGNORANCE the relevant agent's epistemic state is the speaker's, for present purposes the distinction between ascriber and subject is irrelevant. Taking stakes and salience in turn, I will argue that although the presence of factors might be *sufficient* to license the assertion of a conspiracy conditional, it is not plausible that their presence is *necessary* for it. But absent the necessity claim, the error theory is inadequate.

Starting with stakes, it should be clear that an appeal to variations in what *S* has riding on the truth of *p* is a non-starter as part of a general account of the judgment that *S* may sometimes felicitously and canonically assert a conspiracy conditional of the form 'If *p*, *q*'. It may be the case that *p*'s being of sufficient practical concern to *S* licenses an assertion of the conspiracy conditional, but it is not the case that *p* *must* be that way. That is to say, for an ordinary person *S*, there are loads of propositions *p* such that: (i) according to the moderate invariantist, the proposition expressed by 'S knows that *p*' is true in every context; (ii) the question of whether *p* is more or less irrelevant to *S*'s practical interests; and yet (iii) there are contexts in which *S* (or anyone else) may felicitously assert the conspiracy conditional 'If (\neg)*p*, then *q*'. Most of the examples considered in §4 have this feature, but to drive the point home: imagine looking up the weather in any locale you don't care about, and consider the conspiracy conditional:

(42) If it is actually raining there, then the weather report is accurate.

Clearly you may felicitously assert this conspiracy conditional even when you have nothing riding on the weather in that locale. Concerns of stakes, practical interests, etc., are thus orthogonal to the issue of what to make of the felicitous use of conspiracy conditionals. All the worse for error theories that rely on them.

The upshot is that the error-theoretic appeal to *X* will have to lean heavily on the salience condition to explain away conspiracy conditionals. What does this condition amount to? Well, it is not enough that *S* merely *consider* the possibility that $\neg p$. For one, it is easy enough to do that in a context in which the corresponding conspiracy conditional seems *unassertable*. For another, if the exculpatory half of the moderate invariantist's error theory is to be plausible, the fixation has to be the kind of thing that would explain why otherwise rational and competent speakers fail to follow such basic rules about the use of 'knows'. So presumably $\neg p$ is salient to *S* in the relevant way just in case *S* is taking the metaphysical possibility that $\neg p$ *seriously*. What does taking a possibility seriously amount to? Nothing in the literature makes the notion altogether clear. We'll thus have to resort to

taking those who have invoked the mechanism in their theorizing at their word—so see, for instance, the Williamson and Hawthorne passages from above, as well as the famous passage of Lewis’s (1996, p. 549):

Let your paranoid fantasies rip — CIA plots, hallucinogens in the tap water, conspiracies to deceive, old Nick himself — and soon you find that uneliminated possibilities of error are everywhere.

These vague suggestions are enough to get the argument going. Consider again any of the following conspiracy conditionals:

- (1) If Oswald didn’t shoot Kennedy, then someone else did.
- (25) If Federer defeated Murray in straight sets, then this report is accurate.
- (29) If nature is uniform, then induction is a reliable means of forming true beliefs.

If the diagnosis of error in terms of salience is to be plausible, it better be that we find these conditionals natural only when we are in the sorts of psychological states alluded to by Williamson, Hawthorne, and Lewis: those involving paranoid fantasies, the vivid imaginings of metaphysical counterpossibilities, etc. But unfortunately these are *not* the only situations in which conspiracy conditionals sound natural. One’s fantasies can remain unripped even while one thinks or says that if nature is uniform, then induction is a reliable means of forming true beliefs. Likewise, one needn’t have any worries, general or particular, about the reliability of newspapers to think to assert that if Federer defeated Murray in straight sets, then the report is accurate. The point is especially striking when it comes to conspiracy conditionals whose antecedents concerns events in the future, such as:

- (31) If this fair coin lands heads 1,000 times in a row, then most onlookers will be quite confident that the coin is double-headed.³⁶
- (43) If Trump isn’t president tomorrow, then either he’ll have died or resigned.

These conditionals are perfectly assertable even when one isn’t in the business of taking statistical or political anomalies very seriously. There is no non-trivial notion of salience that could explain what distinguishes the cases in which we judge that we know Trump is president from those in which we judge that conditionals like (31) and (43) are felicitous.

But by way of driving the point home, it is worth observing an important contrast between our judgments about skeptical knowledge ascriptions and our judgments about conspiracy conditionals. For most of the mundane propositions the moderate invariantist takes the ordinary person to context-invariantly know—(P₁)-(P₃) continuing to be serviceable examples—it takes some effort to get speakers to start actually *saying* that they aren’t

³⁶ See footnote 24.

known. For example, to get oneself into a setting in which an assertion of ‘I don’t know whether Trump is president of the United States’ seems appropriate, one will probably have to explicitly consider possibilities in which Trump has recently died, or in which some major political event has transpired without the general public knowing, etc. Again, something like this thought is what I take the quote passages from Williamson, Hawthorne, and Lewis to be gesturing at. We should thus expect that when one is *not* in the midst of a fixation of a sufficiently strong sort, then one will *not* be inclined to say things like ‘I don’t know whether Trump is president of the United States’. And this means, in turn, that we should expect *discourse initial* assertions of knowledge denials of (P₁)-(P₃) to generally seem strange. Being discourse initial, these assertions will usually occur in situations in which the negations of (P₁)-(P₃) are not sufficiently salient to the conversational participants.

Is this prediction born out by the facts? Well, imagine a friend walks up to you and asserts one of:

(44) We don’t know whether Trump is president of the United States.

(45) We don’t know whether Oswald shot Kennedy.

Speaking for myself, in the absence of some sort of cue that helps me figure out which possibilities the speaker is worried about, I would find an out of the blue assertion of any of the above jarring. Indeed, if I were to encounter someone who asserted (e.g.) (45), I would be inclined to respond along the lines of: ‘Huh? Do you have some reason to think Oswald *didn’t* shoot Kennedy?’. That is to say, I’d ask for actual evidence—facts about the epistemic states themselves—before I’d accept my interlocutor’s assertion. And until presented with such evidence, I would take myself to *disagree* with them.

But the intuitions are different with conspiracy conditionals. A discourse initial assertion of (1) is a noticeable improvement, as its standing in the philosophical literature suggests. Likewise an assertion of (e.g.):

(22) If Trump is no longer the president of the United States, then it is surprising we haven’t heard about that.

I might still find these speeches a bit puzzling—for instance I might wonder what prompted the person to say such a thing. But I wouldn’t find myself inclined to *disagree* with her, or to find her overly skeptical-minded.

This contrast casts yet further doubt on the use of the mechanisms of salience in an account of our intuitions about conspiracy conditionals. It is implausible that what explains the difference in the discourse initial acceptability of (e.g.) (44) and (22) is that hearing (22) causes me to take possibilities in which Trump is dead sufficiently seriously, while hearing (44) does not.

I thus conclude that the felicitous use of conspiracy conditionals has nothing important to do with facts about salience. The moderate invariantist remains on the hook to explain

why we assert conspiracy conditionals given that IGNORANCE is true. Perhaps she will appeal to some other psychological state—call it *intellectual eagerness*—and claim that its presence causes us to misjudge whether an assertion of a conspiracy conditional ‘If p , q ’ is felicitous. But here the quick complaint about the invariantist’s error theory that we started the section with becomes even more pressing. Whatever intellectual eagerness is, it just doesn’t look like the kind of thing that could explain the way in which ordinary speakers are *in error* when they assert things like (1). That is to say, it makes the exculpatory half of the error theory look beyond repair.

Given IGNORANCE, the moderate invariantist must think that conspiracy conditionals are in every context unassertable on their canonical uses. Ordinary speakers seem not to realize this, and so the moderate invariantist owes us an error theory that explains why. In the absence of a plausible error theory, our judgments about conspiracy conditionals need to be accommodated rather than explained away by our theory of knowledge.

8 Knowledge in the face of conspiracy conditionals

We know from §4 that moderate invariantism about knowledge is not up to the task. This leaves three other options (broadly construed): (i) skeptical invariantism; (ii) subject-sensitive invariantism; and (iii) contextualism. I will conclude the paper with an argument for why contextualism about knowledge is the only real contender here.

Start with skeptical invariantism about knowledge. Being a brand of invariantism, this view has it that if you hold the underlying (epistemic) facts fixed, the truth of ‘ S knows that p ’ in one context c guarantees its truth in every other context c' . But being a skeptical conception of knowledge, it also has it that for an ordinary adult S and boring proposition (P_n), (e.g. (P_1)–(P_3)), ‘ S knows that (P_n)’ is context-invariantly false.

The exact details of the skeptic’s underlying theory of knowledge will not matter for our purposes. So long as it meets the constraint just described, we know the view will have a straightforward explanation of why conspiracy conditionals are sometimes felicitous. The issue, however, is that we also know that the view will suffer from the defect §6.2 argued the “certainty-theoretic” conception of OPENNESS suffers from. The defect is that it cannot offer a principled explanation of: (i) the fact that sometimes we may felicitously assert sentences expressing propositions like (P_1) (that Oswald shot Kennedy), (ii) the fact that sometimes we may felicitously assert the corresponding conspiracy conditionals, but (iii) we can never felicitously assert the conjunction of these sentences with their corresponding conspiracy conditionals. In short: because by the skeptic’s lights we know so little, it’ll always be very easy to satisfy the demands of IGNORANCE. Why, then, is it also so easy for conspiracy conditionals to seem *infelicitous*? Absent a concrete answer to this question, the view fares no better than its moderate counterpart.

Next there is sensitive invariantism about knowledge.³⁷ The sensitive invariantist is a moderate invariantist who takes the facts about what's at stake for, or salient to, the subject of a knowledge ascription to context-invariantly affect the ascription's truth-conditions. Put another way, the sensitive invariantist takes the X-like features of the moderate invariantist's error theory and plugs them into the semantics for 'knows' itself. What this means is that whether one knows a proposition depends not just on one's beliefs and evidence concerning that proposition; it also depends on what one has at stake on the truth of the proposition, and/or on whether one finds salient possibilities in which the proposition is false. Thus, you can have two agents *S* and *S'* for whom the purely epistemic facts (beliefs, evidence, etc.) are identical as regards *p*, but for whom the facts about stakes and salience are *not* identical as regards *p*, and as a result 'S knows that *p*' will be context-invariantly true while 'S' knows that *p*' will be context-invariantly false.

Sensitive invariantism has been thought to capture much of the data that motivates contextualism about knowledge, while also avoiding some of contextualism's more exotic meta-semantic commitments.³⁸ However, if the arguments just given against the moderate invariantist's error theory of conspiracy conditionals are sound, then they should tell more or less equally against the sensitive invariantist's semantics for 'knows'. To account for conspiracy conditionals like (1), the sensitive invariantist would have to argue that the psychological factors that make (1) seem assertable have the effect of destroying one's knowledge that Oswald shot Kennedy—no matter one's beliefs and evidence on the matter. This treatment of conspiracy conditionals may be even less plausible than the error-theoretic one offered by the moderate invariantist.

It looks like the facts about the felicitous use of conspiracy conditionals tell against all non-contextualist theories of knowledge alike. In any given context, there is an intimate connection between judgments about the truth of 'S knows whether *p*' and the truth of 'S may assert 'If *p*, *q*'. In particular, the truth of the former implies the falsity of the latter. But even holding the underlying epistemic facts fixed, there is clear variation across contexts in the truth of 'S may assert 'If *p*, *q*'. The differences between these contexts do not look to be of a kind amenable to explanations error-theoretic or subject-sensitive.

This leaves us with the contextualist's theory of 'knows'. The contextualist accepts that in any context in which the proposition expressed by 'S may assert 'If *p*, *q*' is true, the proposition expressed by 'S knows whether *p*' is false. Her story about the felicitous use of conspiracy conditionals is that these uses arise in contexts in which, by the lights of the context, 'S knows whether *p*' is false. Simple as that. Of course, there remains the question of the mechanism that explains the contextual variability in the appropriateness

³⁷ As I will use the label 'sensitive invariantism', one who believes knowledge is subject to "pragmatic encroachment" (and isn't a contextualist) is a sensitive invariantist. See Kim (2017) for a helpful overview.

³⁸ See, e.g., issues related to "semantic blindness" discussed by Schiffer (1996); Williamson (2005); Greenough & Kindermann (2017).

of assertions of conspiracy conditionals. But that's a question about the contextualist's *meta*-semantics, not her semantics. One need not have a theory of the mechanisms driving the context-sensitivity of an expression to know that it is context-sensitive.³⁹ And in the present case we know enough about the norms governing the felicitous use of indicative conditionals and the actual facts of their usage to know that the correct semantics of 'knows' must be the contextualist's. Thus, conspiracy conditionals do more than just break the stalemate between various conceptions of knowledge; they constitute independent and deceptively strong grounds in favor of contextualism.

³⁹ [Ichikawa \(2011\)](#) makes a similar point in defense of [Lewis's \(1996\)](#) theory of 'knows'.

References

- Adams, Ernest W. 1970. Subjunctive and Indicative Conditionals. *Foundations of Language*, 6(1), 89–94.
- Adler, Jonathan E. 2006. Withdrawal and Contextualism. *Analysis*, 66(4), 280–285.
- Austin, J. L. 1946. Other Minds. *Supplementary Proceedings of the Aristotelian Society*, 20, 149–187.
- Bach, Kent. 2005. The Emperor's New 'Knows'. *Pages 51–89 of: Preyer, Gerhard, & Peter, Georg (eds), Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford University Press.
- Bacon, Andrew. 2014. Giving Your Knowledge Half a Chance. *Philosophical Studies*, 1–25.
- Buckwalter, Wesley, & Schaffer, Jonathan. 2015. Knowledge, Stakes, and Mistakes. *Noûs*, 49(2), 201–234.
- Cohen, Stewart. 1986. Knowledge and Context. *Journal of Philosophy*, 83(10), 574–583.
- Cohen, Stewart. 1999. Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives*, 13(s13), 57–89.
- Declerck, Renaat, & Reed, Susan. 2001. *Conditionals: A Comprehensive Empirical Analysis*. Berlin: Mouton de Gruyter.
- DeRose, Keith. 1992. Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, 52(4), 913–929.
- DeRose, Keith. 1995. Solving the Skeptical Problem. *Philosophical Review*, 104(1), 1–52.
- DeRose, Keith. 2002. Assertion, Knowledge, and Context. *Philosophical Review*, 111(2), 167–203.
- DeRose, Keith, & Grandy, Richard E. 1999. Conditional Assertions and "Biscuit" Conditionals. *Noûs*, 33(3), 405–420.
- Dorr, Cian, & Hawthorne, John. 2013. Embedding Epistemic Modals. *Mind*, 122(488), 867–914.
- Dorr, Cian, Goodman, Jeremy, & Hawthorne, John. 2014. Knowing Against the Odds. *Philosophical Studies*, 170(2), 277–287.
- Edgington, Dorothy. 1995. On Conditionals. *Mind*, 104(414), 235–329.
- Edgington, Dorothy. 2014. Indicative Conditionals. *In: Zalta, Edward N. (ed), The Stanford Encyclopedia of Philosophy*, winter 2014 edn. Metaphysics Research Lab, Stanford University.
- Fantl, Jeremy, & McGrath, Matthew. 2002. Evidence, Pragmatics, and Justification. *Philosophical Review*, 111(1), 67–94.
- Fantl, Jeremy, & McGrath, Matthew. 2009. *Knowledge in an Uncertain World*. Oxford University Press.

- Gazdar, Gerald. 1979. *Pragmatics: Implicature, Presupposition and Logical Form*. New York: Academic Press.
- Gillies, Anthony. 2010. Iffiness. *Semantics and Pragmatics*, **3**(4), 1–42.
- Gillies, Anthony S. 2004. Epistemic Conditionals and Conditional Epistemics. *Noûs*, **38**(4), 585–616.
- Greenough, Patrick, & Kindermann, Dirk. 2017. The Semantic Error Problem for Epistemic Contextualism. *Pages 305–320 of: Ichikawa, Jonathan (ed), Routledge Handbook of Epistemic Contextualism*. Routledge.
- Grice, H. Paul. 1975. Logic and Conversation. *Page 47 of: Ezcurdia, Maite, & Stainton, Robert J. (eds), The Semantics-Pragmatics Boundary in Philosophy*. Broadview Press.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford University Press.
- Ichikawa, Jonathan. 2011. Quantifiers and Epistemic Contextualism. *Philosophical Studies*, **155**(3), 383–398.
- Kaplan, David. 1989. Demonstratives. *Pages 481–563 of: Almog, Joseph, Perry, John, & Wettstein, Howard (eds), Themes From Kaplan*. Oxford University Press.
- Khoo, Justin. 2015. On Indicative And Subjunctive Conditionals. *Philosophers' Imprint*, **15**.
- Kim, Brian. 2017. Pragmatic Encroachment in Epistemology. *Philosophy Compass*, **12**(5), e12415.
- King, Jeffrey C., & Lewis, Karen S. 2017. Anaphora. *In: Zalta, Edward N. (ed), The Stanford Encyclopedia of Philosophy*, summer 2017 edn. Metaphysics Research Lab, Stanford University.
- Leahy, Brian. 2011. Presuppositions and Antipresuppositions in Conditionals. *Proceedings of Semantics and Linguistic Theory*, 257–274.
- Lewis, David. 1996. Elusive Knowledge. *Australasian Journal of Philosophy*, **74**(4), 549–567.
- May, Joshua, Sinnott-Armstrong, Walter, Hull, Jay G., & Zimmerman, Aaron. 2010. Practical Interests, Relevant Alternatives, and Knowledge Attributions: An Empirical Study. *Review of Philosophy and Psychology*, **1**(2), 265–273.
- McGrath, Matthew. 2007. On Pragmatic Encroachment in Epistemology. *Philosophy and Phenomenological Research*, **75**(3), 558–589.
- Nagel, Jennifer. 2008. Knowledge Ascriptions and the Psychological Consequences of Changing Stakes. *Australasian Journal of Philosophy*, **86**(2), 279–294.
- Nagel, Jennifer. 2010. Knowledge Ascriptions and the Psychological Consequences of Thinking About Error. *Philosophical Quarterly*, **60**(239), 286–306.
- Predelli, Stefano. 2009. Towards a Semantics for Biscuit Conditionals. *Philosophical Studies*, **142**(3), 293–305.
- Reed, Baron. 2010. A Defense of Stable Invariantism. *Noûs*, **44**(2), 224–244.

- Roeber, Blake. 2018. The Pragmatic Encroachment Debate. *Noûs*, 171–195.
- Rothschild, Daniel. 2012. *A Note on Conditionals and Restrictors*.
- Rysiew, Patrick. 2001. The Context-Sensitivity of Knowledge Attributions. *Noûs*, **35**(4), 477–514.
- Schaffer, Jonathan. 2005a. The Irrelevance of the Subject: Against Subject-Sensitive Invariantism. *Philosophical Studies*, **127**(1), 87–107.
- Schaffer, Jonathan. 2005b. What Shifts? : Thresholds, Standards, or Alternatives? *In*: Preyer, Gerhard, & Peter, Georg (eds), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford University Press.
- Schaffer, Jonathan. 2007. Knowing the Answer. *Philosophy and Phenomenological Research*, **75**(2), 383–403.
- Schaffer, Jonathan, & Szabo, Zoltan Gendler. 2013. Epistemic Comparativism: A Contextualist Semantics for Knowledge Ascriptions. *Philosophical Studies*, 1–53.
- Schiffer, Stephen. 1996. Contextualist Solutions to Scepticism. *Proceedings of the Aristotelian Society*, **96**(1), 317–333.
- Stalnaker, Robert. 1975. Indicative Conditionals. *Philosophia*, **5**(3), 269–286.
- Stalnaker, Robert. 2002. Common Ground. *Linguistics and Philosophy*, **25**(5-6), 701–721.
- Stalnaker, Robert C. 1999. *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford University Press UK.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford University Press.
- Unger, Peter K. 1975. *Ignorance: A Case for Scepticism*. Oxford University Press.
- von Fintel, Kai. 1999. The Presupposition of Subjunctive Conditionals. *Pages 29–44 of*: Sauerland, Uli, & Percus, Orin (eds), *The Interpretive Tract*. MIT Working Papers in Linguistics 25.
- Weatherson, Brian. 2012. Knowledge, Bets, and Interests. *Pages 75–103 of*: Brown, Jessica, & Gerken, Mikkel (eds), *Knowledge Ascriptions*. Oxford University Press.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.
- Williamson, Timothy. 2005. Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge. *The Philosophical Quarterly*, **55**(219), 213–235.
- Worsnip, Alex. 2017. Contextualism and Knowledge Norms. *In*: Ichikawa, Jonathan (ed), *The Routledge Handbook of Epistemic Contextualism*. Routledge.