

Infallibilism and Easy Counter-examples

Alex Davies (forthcoming in *Grazer Philosophische Studien*)

Abstract: Infallibilism is commonly rejected because it is apparently subject to easy counter-examples. I describe a strategy that infallibilists can use to resist this objection. Because the sentences used in the counter-examples to express evidence and belief are context-sensitive, the infallibilist can insist that such counter-examples trade on a vacillation between different readings of these sentences. I describe what difficulties await those who try to produce counter-examples against which the proposed strategy is ineffective.

1 The example based objection

Infallibilism about evidence-based belief is the following thesis:

Infallibilism

A knows that P on the basis of evidence E only if E entails P.¹

Infallibilism is the kind of philosophical thesis that undergraduates have to be taught out of believing—in the first instance, untutored intuition tends to be in its favour.

But there is an objection that is commonly raised against infallibilism. It looks like an easy objection to make and it seems to render infallibilism implausible.

Example based objection

(1) We have many evidence-based beliefs that we intuitively take to be knowledge but which are not entailed by the evidence they're based upon.

(2) If (1) then infallibilism implies scepticism.

Therefore:

(3) Infallibilism implies scepticism.

There are various ways one might defend (1). But the version of this argument that I call “the example based objection” incorporates the following defence of (1): purportedly unequivocal examples of evidence-based beliefs which are not entailed by the evidence upon which they are based, but which do intuitively constitute knowledge, are described.² This use of examples to defend (1) is deployed in many places e.g. (Cohen, 1988), (Dodd, 2007), (Dretske, 1970, pp. 1015-1016), (Dutant, 2007), (Fantl & McGrath, 2009, p. 59), (Lewis,

¹ I restrict attention to evidence-based belief because I don't want infallibilism to turn out false simply because there's such a thing as knowledge that is not based on evidence. Defenders of infallibilism (more or less so understood) include Bonjour (2010), Dodd (2011), Leddington (forthcoming), McDowell (1982), and Travis (2005). Unlike the others, Bonjour and Dodd both defend infallibilism *and* either argue (in the case of (Dodd, 2007)) or grant (in the case of (Bonjour, *ibid*)) that infallibilism implies scepticism.

² In particular, I am not interested in defences of (1) that attempt to derive it from further assumptions about the nature of evidence or the ways in which it is acquired. For instance, one might first show that our evidence consists of our non-factive mental states and then use that conclusion to derive (1). I will not address such defences here.

1996), (Littlejohn, 2008), (O'Brien, 2006, p. 15), and (Reed, 2012, p. 586)).³ If the argument is sound, infallibilism is a very controversial thesis. Most philosophers (including most of those listed just now) think scepticism is false. Indeed, this is the main reason why fallibilism is the preferred thesis in this domain.⁴

In this paper I identify a criticism that the example based objection can face. I think that basically all extant attempts to deploy the objection are subject to this criticism. I do not mean to suggest that this criticism shows that the example based objection cannot be used against infallibilism. Rather, I aim to show that compelling applications of the example based objection require more argumentative work than they might at first seem to need. Consequently, the example based objection is not the *easy* route to rejecting infallibilism that it is so often treated as being.

Criticism of the objection is warranted when the examples provided in support of (1) are not unequivocal examples with the two required features: evidence-based beliefs that are both not entailed by the evidence upon which they're based but which intuitively do constitute knowledge. If those examples use context-sensitive sentences to express either the evidence itself or the beliefs that are based on that evidence, then, no doubt, these sentences can express propositions that do not stand in a relation of entailment. But they can also express subtly different propositions that do stand in a relation of entailment. It is therefore possible that the examples trade on a vacillation between two readings of the sentences used to express the beliefs and evidence. When we acknowledge that the belief is something we take ourselves to know, that is because the sentences are understood to express propositions that stand in a relation of entailment. But when we acknowledge that the belief is not entailed by the evidence upon which it is based, that is because the sentences are understood to express propositions that do not stand in a relation of

³ Three comments are in order. Firstly, Dodd (2007) defines infallibilism as the thesis that knowledge that *p* requires that the probability that *p*, conditional upon one's evidence, is 1. Dodd (2011) defines infallibilism as the thesis that *S*'s knowledge that *p* requires that the epistemic probability that *p* for *S* be 1. These definitions diverge from the definition used in this paper. The former definition may come apart from our own in cases where the probability of *p* conditional upon one's evidence is 1, but one's evidence doesn't entail *p*. The latter definition may come apart from ours in cases described by Dougherty (2011, pp. 140–141). For our purposes, the differences don't matter—because our concern is with the fact that Dodd (2007) is trying to support (1) using examples. So I ignore such differences in this paper. Secondly, in case you're wondering why Lewis is in this list, Lewis' contextualism is not infallibilism. See (Doutven, 2005). Thirdly, although Dutant (2007, 2016) defends a thesis that he calls "infallibilism," it is not the thesis discussed here—which he rejects.

⁴ There is one rather popular response to the example based objection, which deserves mention. Several philosophers accept the argument (1)-(3) as sound, but deny that scepticism is a counter-intuitive thesis (e.g. (Bonjour, 2010, pp. 74–76), (Davis, 2007, *passim*) and (Fumerton, 2010, pp. 248–251)). They argue that in everyday talk we use "know" loosely. So although, as the soundness of the argument would imply, most of our knowledge-claims are strictly speaking false, we are making those claims in order to communicate something that is true (e.g. that one's belief is close enough to knowledge for practical purposes). Dodd, on the other hand, simply grants that infallibilism implies scepticism, and defends infallibilism—suggesting a kind of nonchalance on his part about commitment to scepticism. The response I pursue in this paper provides the infallibilist with a strategy for establishing the argument's unsoundness and thus for avoiding the concession that infallibilism implies scepticism.

entailment. This possibility needs to be convincingly eliminated from an example presented in support of (1) in order for the objection against infallibilism to go through.

To illustrate the criticism, in sections 2-4 I apply it to three attempts (by Dodd (2007) and Dutant (2007)) to support (1) using examples. The examples they provide concern beliefs about the immediate past, the not immediately observed, and the future. But each example uses context-sensitive sentences to express the relevant beliefs and evidence and these sentences do have readings on which the evidence entails the beliefs based upon it. No attempt is made to convincingly eliminate the possibility that the examples only appear to have the status of knowledge without entailment because of a vacillation in how we read the relevant sentences. It is therefore open to an infallibilist to insist that the examples trade on such a vacillation and do not constitute phenomena to be saved by an adequate theory of knowledge. In section 5, I describe two problems that face attempts to produce examples in support of premise (1) which are not susceptible to the criticism I describe. In section 6, I explain why attempts to produce examples in support of premise (1) using inductive beliefs are just as susceptible to the criticism as beliefs about the future, the past and the not immediately observed.

Two last things before we begin. Firstly, I should be up-front about the fact that I will be taking evidence in favour of context-sensitivity at face value. I assume that when it appears that a sentence can have different truth-values when used in different contexts, to speak of the same object, in the same state, then the sentence's content is different across the two contexts. One might reject this assumption on two different grounds. Firstly, one might reject it because one simply doesn't share the judgements I (and my informants, and those I cite) have on these matters. Secondly, one might share the judgements but insist that they are to be understood as tracking propositions that a sentence pragmatically implicates in a context, and not context-sensitive semantic contents of the relevant sentences. In response to the first ground for rejection, I don't have anything to say. I assume that the readers of this paper mostly share the judgements of myself and my informants and those I cite on these matters. But in response to the second ground for rejection, I want to stress that it is not a way to escape the conclusion of this paper. The reason being that if one grants the judgements that I make in this paper (i.e. one doesn't pursue the first ground for rejection), then one has to grant that thinkers using the sentences we discuss to express their evidence and evidence-based beliefs will take their evidence and evidence-based beliefs to be what is pragmatically implicated by, and not the semantic content of, these sentences. For it's those propositions which are most saliently expressed by the use of the sentences we discuss. Hence, assuming that we don't suffer from a general ignorance of what our evidence and evidence-based beliefs are, when we express them, it must be these propositions that are our evidence and evidence-based beliefs. But because the propositions pragmatically implicated by these sentences nonetheless vary with context, the potential for illicit vacillation between two different propositions in examples given in support of premise (1) will remain.

Secondly, Dodd himself defends infallibilism in one paper, while arguing that it implies

scepticism in another. In neither paper does he reject infallibilism on the ground that it implies scepticism. This might make one doubt that the examples of his that we will discuss are really intended to support (1). For, given his commitments, surely he doesn't think that there are any examples of evidence-based beliefs that constitute knowledge even though they are not entailed by the evidence upon which they are based. However, although he is committed to denying the existence of such beliefs, nonetheless, (1) is about what beliefs intuitively seem to be knowledge: it is not about which beliefs actually constitute knowledge. Dodd has to be in the business of providing examples of beliefs that appear to be knowledge if he wants to conclude that infallibilism implies scepticism. For the mere fact that we have a great many beliefs that don't constitute knowledge, doesn't constitute or imply scepticism. Everyone (more or less) agrees that many beliefs formed in a paranoid rage, or whilst drunk, etc. don't constitute knowledge. And there are a great many of those. But pointing this out, isn't a way of supporting scepticism. Scepticism is a thesis that conflicts with intuition, stating that things we took ourselves to know, we don't in fact. Thus Dodd must be in the business of providing examples in support of (1), even if he defends a view he takes to imply scepticism.

2 Dutant on the immediate past

Because observations made in the discussion of his example will also be useful when discussing one of Dodd's examples, let's begin with Dutant (2007). He presents the following example:

[T]here are intuitions on standard knowledge cases. Suppose one believes that there was a coin on the table five seconds ago on the basis of one's knowledge that there is a coin on the table now. What one knows is consistent with the falsity of that belief. However, intuitively, it seems possible for one to come to know that there was a coin on the table five seconds ago in such a way, at least in standard situations. (*ibid*, p. 72)

Dutant thinks he has an unequivocal example of a belief with two features: we take it to be possible to know it on the basis of the evidence provided but the belief is not entailed by that evidence. Dutant supposes that the proposition expressed by (4) doesn't entail the proposition expressed by (5):

- (4) There is a coin on the table now.
- (5) There was a coin on the table five seconds ago.

For a moment, set aside the temporal adverbials and focus on the difference in tense. So we have:

- (4') There is a coin on the table.
- (5') There was a coin on the table.

Does the proposition expressed by (4') entail the proposition expressed by (5')? It *can* do so.

To see this, we need to take note of two features of tense in English.

Firstly, we should not understand past tense claims as claims to the effect that, at some time or other prior to the time of utterance, such and such happened or was happening. Consider the following sentence:

(6) I left the stove on.

For typical uses of a sentence like (6), it is not the case that the sentence is true only if there is some time or other prior to the time of utterance at which the speaker left the stove on, but rather, only if the speaker left the stove on at some particular time prior to the time of utterance—a time which is fixed by the context of utterance (cf. (Partee, 1973)). Past tense sentences ought often to be understood as requiring that something be so of some particular time in the past and not just of some time or other in the past.

Secondly, the times that tensed sentences are about can be *instants* of time or *intervals* of time. Tensed sentences are often and perhaps normally about intervals of time (cf. (Bennett & Partee, 1972)). So, if (6) is true, there's not just an instant of time—a flash of a second—during which the speaker left the stove on but, more likely, some contextually fixed interval during which the stove was not turned off—e.g. during the minutes in which the speaker was leaving the house, or during the time when he was finishing cooking, or something similar.

Both of these features are as much features of the present tense as they are of the past: present tense sentences are typically about an interval of time and which interval that is fixed by the context of use (for discussion see (Klein, 1994)). For example, compare the following two sentences. The first is about a pair of race horses said during a race commentary and the second is about a famous building in London said whilst giving directions to someone:

(7) Seldom Seen is just to the front of Big Ben.

(8) The Thames is just to the front of Big Ben.

These two sentences, as used in these contexts, differ significantly: there's no commitment that things will stay like they are described in (7) for even the time taken to say another sentence (thus the need for a racing commentator to speak rather quickly). But the period of time for which the Thames is thus related to Big Ben is significantly longer (which is why you can use the directions to find your way to the Thames). Moreover, sometimes the present can begin at a time prior to the onset of the utterance time. Consider:

A is talking to himself while examining a crime scene wherein the culprits did break the window but did not enter:

(9) A: I'm wondering why they wanted to get in. I guess, there *are* two coins on the table.

A bird flies around the room:

(10) A: How did it get in here? B: The window's open.

(11) A: Why is it so warm in here? B: The thermostat is set to 25 degrees.

(12) A: Why was your brother so angry? B: I'm wearing his jeans.

In each case, we naturally take the relevant present tense claim to be about a stretch of time that extends into the past beyond the onset of the utterance of the relevant sentence so that insofar as the relevant state of affairs does not extend into the past in this way, the statement will not be accepted as true.

Now, if the past tense sentence (5') can be about a contextually indicated interval of time and if (4') can be about an interval of time that extends into the past prior to the onset of the utterance of (4'), then, it is possible that (i.e. there will be some contexts of use of (4') and (5'), such that) the interval of time that (4') is about includes the interval of time that (5') is about. If that's the case, then it's quite possible for (4') to express a proposition that entails the proposition expressed by (5').

This doesn't necessarily change when we add in the temporal adverbials. Let's return to (4) and (5):

(4) There is a coin on the table now.

(5) There was a coin on the table five seconds ago.

Consider what propositions (4) and (5) can express. I assume that the temporal adverbial in (5) is fairly precise in how it constrains the time that the sentence is about: it makes (5) about an interval of time that begins approximately five seconds prior to the time of utterance and extends further back into the past by some contextually fixed extent. But the temporal adverbial in (4) is wildly context sensitive. Assuming that "now" is about the present, it'll be about an interval and which interval it is about will vary with context—even when the time of utterance is held constant, the duration and temporal location of the beginning and end of that interval can vary with context. Moreover, that interval can extend back into the past beyond the onset of the utterance, as we see in (13)-(16):

(13) Things change with the centuries. In the 1500s, women were accused of being witches. This does not happen now.

(14) The economic situation is now much better than it was in the *previous* decade.

(15) The library is not open during the summer but it's open now. It's autumn.

(16) Speaking to someone who visited the exhibition only last year: we set up the exhibition a little differently this year. Last year, there was a brooch on the table but there is a coin on the table now.

These observations strongly suggest that it is possible for the interval of time that (4) is about to include the interval of time that (5) is about—i.e. there are some contexts of use of

(4) and (5) such that this is so. But then, in such contexts, (4) expresses a proposition that entails the proposition expressed by (5). If so, then, even given infallibilism, it's still possible for one to know the proposition expressed by (5) when the proposition expressed by (4) is one's evidence. For example:

Exhibition Context

Dodd and Dave are discussing how this year's exhibition differs from last year's exhibition. Last year there was a dead cat on the table. Dave explains to Dodd how this year's exhibition is laid out—an exhibition which started 3 days ago:

Dave: Last year there was a dead cat on the table. But there is a coin on the table now.

In Exhibition Context, “now” refers to the period of the exhibition and thus extends back three days. Thus the proposition Dave asserts, in this context, using the sentence “there is a coin on the table now” entails the proposition that there was a coin on the table five seconds prior to the time of utterance.

With all this in mind, does Dutant's example have the two required features? We cannot say that it unequivocally does. He's right that 'intuitively, it seems possible for one to come to know that there was a coin on the table five seconds ago' on the basis of 'one's knowledge that there is a coin on the table now.' But that intuition could be being elicited by the fact that the sentence 'there is a coin on the table now' can express a proposition that entails that there was a coin on the table five seconds ago. In order for Dutant's example to have the two features unequivocally, we need some assurance that our intuition does not arise from this source. Dutant provides no such assurance.

Of course, we could change the example in such a way that the particular manifestations of context-sensitivity described in this section can no longer allow for a relation of entailment between evidence and belief. The same will be so for the examples discussed in sections 3 and 4. But in each case, this would only raise the question afresh of whether context-sensitivity is present which can be used by an infallibilist to make a parallel criticism of the (new) example. The trick that's needed is to produce an example where some such context-sensitivity is not available. We'll return to this issue in section 5.

3 Dodd on the not immediately observed

Dodd (2007) thinks that many beliefs we have about the present are beliefs that we take ourselves to know but which are not entailed by the evidence upon which they're based. He has in mind beliefs we have about the not immediately observed. He defends this contention by providing an example of a belief about the not immediately observed which we would take to be knowledge but which is not entailed by the evidence upon which it is based. He begins by telling us a story about how a desk can cease to be a desk:

Suppose that there is a desk in front of me. Quantum mechanics tells us that there is

a wave function that describes the space of nomically possible developments of the system that is my desk. On those interpretations of quantum mechanics according to which the wave function gives probability of location, there is some non-zero probability that, within a short while, the particles belonging to the surface of the desk remain more or less unmoved but the material inside the desk unfolds in a bizarre enough way that the system no longer counts as a desk. (Dodd (2007, p. 646) quoting from (Hawthorne, 2004, pp. 4–5))

Given an externalism about the content of what one perceives, Dodd grants that when Dodd is looking at his desk in his office, he may have evidence that entails that there is a desk in his office. Moreover, Dodd grants that, when he loses perceptual contact with the desk, he can still have evidence that entails that there *was* a desk in his office because his memory's content is fixed by the content of his original perceptual state:

I may remember that there was a desk in my office a short while ago, because the content of my perceptual state of remembering that there was a desk in my office is determined in part by the desk itself (or perhaps a past stage of the desk itself).(Dodd, 2007, p. 648)

However, Dodd denies that, at the later time, he has evidence that entails that there *is*, at that time, a desk in his office—because his evidence is consistent with the possibility that the desk has spontaneously transformed into a non-desk in the meantime:

But right now I am not perceiving that there is presently a desk in my office. I believe that there is presently a desk in my office because I remember that there was a desk there a short time ago, and even though I am unable to observe that there is a desk there right now, and even though I am unable to observe that there is a desk there right now, I have no reason to think that it is not there still... I believe that there is a desk there only because I remember that it was there a short time ago, and because of assumptions I am making about how desks tend to stay in existence and in the same location. These memories and assumptions provide me with the grounds or evidence for my belief that there is a desk in my office right now. They give me all the information I have to go on. (*ibid*, pp. 647-648)

From this, Dodd (*ibid*, p. 649) concludes that, if knowledge requires possession of evidence that entails the truth of what one believes on its basis, then:

I can know almost nothing about what sorts of objects currently exist outside the small region of the world I am able to observe directly

However, it is quite possible that the content of Dodd's memory entails that there *is* a desk in Dodd's office. Consider the sentence:

(17) There is a desk in my office.

This is a sentence in the present tense. Sentences in the present tense are about the present: they require that some state of affairs or process obtains during the present. However, the present is best understood as an interval and not a moment and which interval this is will depend on the context in which the present tense sentence is employed—again, even with the time of utterance held fixed. Compare the following:

Could you tell me how your office is arranged so that I know what to expect when I arrive in several weeks time?

- (18) There is one door. There is a window. There is a desk in my office. There may be a cat.

During the removal of furniture from the office, because Dodd is moving to a new job, A asks “Is there anything left?” and Dodd replies:

- (19) There's a desk in my office.

In (18), at least on one salient reading, in order for the proposition expressed by “there is a desk in my office” to be true, the desk must be there during the time between the time of utterance and the time when the person asking the question arrives in the office. If this isn't so then there's a familiar kind of presupposition failure. It would be like saying that there is a chair when a chair is cyclically disappearing and reappearing (cf. Wittgenstein's (2001) discussion of a disappearing chair at paragraph 80, and Bennett and Partee's (1972) sub-interval property). In (19) the interval during which the desk must be in the office in order for the sentence to be true is much briefer.

So the period of time during which there must be a desk in Dodd's office in order for the sentence “there is a desk in my office” to be true, depends on the context for uttering the sentence. In principle, then, that interval of time could extend right through to the time at which Dodd is remembering the condition of his office. This possibility is reflected in the fact that there are two ways that Dodd could report what he is remembering:

- (20) I remember that there was a desk in my office.

- (21) I remember that there is a desk in my office.

(21) is felicitous and can be true. Imagine, for instance, that the speaker of (21) has been in an accident, and she is asked what she remembers about her workplace. She gives (21) as part of her answer. Or suppose that two people are on their way to stay at a friend's apartment and one (who hasn't been at the apartment before) asks the other (who has) whether she remembers whether there's a place to write in this apartment. The response given is a variation on (21): “Well, I remember that there is a desk in her office. Though I'm not sure she'll let us use it.” As the felicity and possible truth of (21) shows, how things are in the present is something one can remember. Moreover, as we have seen, given the context-sensitivity of the size of the present interval, how things are at the time of utterance is a possible part of the content of Dodd's original perception. But then, if one

can perceive that there is a desk in one's office and if this is a state of affairs that extends into the future (because the contextually specified present interval is rather long), and, finally, if one can have a memory with this content, then what one remembers can entail that there is (currently) a desk in one's office. But then Dodd is wrong when he says that the memory that there *was* a table in his office a short while ago and his assumptions about the persistence of objects 'give me all the information I have to go on' (*ibid*, p. 648).

With all this in mind, does Dodd's example have the two required features? While we cannot say that it definitely does not, we also cannot say that it unequivocally does. He's right that one can come to know that there is a desk in one's office on the basis of the evidence one acquired by seeing it there at some point in the past. But he's wrong to suppose that this evidence obviously doesn't entail the truth of his belief. In order for Dodd's example to have the two features unequivocally, we need some assurance that the reason why we take it to be possible that we can know that there is a desk in one's office on the basis of the relevant perceptual evidence is not that the evidence entails the truth of the belief (as it can do so). Dodd provides no such assurance.

At this point, one might acknowledge the felicity of (21) but refuse point-blank to grant that one can remember how things *are* and insist that one can only remember how things *were*. One might insist that when one utters things like (21), one says things that are strictly speaking false, even if something else is implicated which is true (*viz.* that there is something that was the case, that one remembers).⁵ But even if one were right to so insist, it would still seem to us as though we can know things about the present based on what we can remember about the present. So it would still seem to us that part of our evidence is what we can remember about the present and that this entails what we believe on its basis. It would then still be possible that the sentences used to express this (even though, in part, strictly speaking false) seem to us to have two readings: one on which what we take to be our evidence entails what we believe on its basis, and one on which it does not. Thus assurance would still be required that the examples of this kind provided in support of (1) involve no illicit vacillation of (apparent) readings.

4 Dodd on the future

Dodd (2007) claims to provide two examples of beliefs about the future which we take to be knowledge but which are not entailed by the evidence upon which they're based.⁶

Dodd's (2007, pp. 644-64) first example is a belief expressed using the following sentence:

(22) I shall not be able to afford an African safari.

⁵ Note that the envisaged concern raised here about pragmatic content is different from that considered in the introduction to this paper. In the introduction, I considered the possibility that a reader may think it better to interpret the examples I present as eliciting intuitions that track pragmatic, and not (context-sensitive) semantic, content. But the concern envisaged at this stage of the paper is that although our intuitions do track semantic content, they do not accurately track the truth-value of this content.

⁶ What is said here, can also be said of Littlejohn's (2008) variation on Dodd's examples.

Dodd's beliefs about 'what is going to be the case in the future' are based on evidence that he has now (Dodd, 2007, p. 645). Dodd assumes that the evidence he has now includes the fact that he has bought a lottery ticket. He thinks that conditional upon this evidence, there is a non-zero chance that he will win the lottery. But, he thinks, if he wins the lottery then the belief he expresses with (22) is false. Thus, his belief expressed with (22) is something which would seem, often enough, to constitute knowledge but which is not entailed by his evidence.

Dodd's (2007, p. 645) second example is any belief, F , about what he will be doing in the future. Dodd supposes that his evidence is consistent with the possibility that he dies before doing the things he believes he will be doing in the future. He also supposes that his beliefs about what he will be doing in the future are false if he dies before doing them. Thus such beliefs—given that they often enough seem to constitute knowledge—support premise (1) of the example based objection.

I'll now argue that beliefs about the future are expressed using context-sensitive sentences—and this means both that the belief Dodd expresses with (22) can be true even if he wins the lottery and that his belief about what he will be doing can be true even if he dies before doing them.

English has no verb form that corresponds to talk about the future—in this sense, English has no future tense. Nonetheless, there are several devices that can be used to make one's talk about the future. We will focus on two (both used by Dodd in support of (1)). Firstly, one can construct a sentence with the modal verb “will” as in “John will run” or the somewhat archaic and formal “shall” as in “John shall run.” In what follows, I will treat these sentences as interchangeable. I will call them “w-sentences.” Secondly, one can construct a sentence with the expression “be going to” as in “John is going to run.” I will call these “bgt-sentences.” Beliefs about the future expressed using w-sentences of the form “A will ϕ ” and bgt-sentences of the form “A is going to ϕ ” can be true even though A does not ϕ in the future.

Let's begin with bgt-sentences. Consider the difference between the following two sentences:

- (23) The Queen was going to pass along that road later but she did not.
- (24) The Queen would pass along that road later but she did not. (Palmer, 1979, p. 149)

(24) appears contradictory but (23) does not. This suggests that beliefs about what is going to happen are compatible with it not actually happening. It can be true that A was going to ϕ even if A never goes on to ϕ . Similarly, an utterance of “A is going to ϕ ” (i.e. in the present tense) can be retrospectively, correctly assessed as true, even if A does not go on to ϕ . Thus, it can be true that the tree is going to fall, even though, in the end, Superman swoops in to stop it from falling. It can be true that the bleach is going to ruin the clothes,

even though your brother saves the day by rescuing the clothes from the washing machine just in time. It can be true that the town is going to flood even though at the last moment the government blows up a wall to safely redirect the flood water.

This suggests that bgt-sentences are made true by something about the present and not by how things actually turn out. In the words of Palmer:

BE GOING TO is used to suggest, in its present tense forms, that there are features of the present time that will determine future events. It is thus essentially a marker of 'future in the present'. (Palmer, 1979, p. 144)

This fact about bgt-sentences is consequential. Dodd thinks that if, conditional upon his evidence, there is some small chance that he is not alive next year then, conditional upon his evidence, there's some small chance that his beliefs about what he is going to be doing next year are not true and that consequently his evidence does not entail his beliefs about what he is going to be doing next year. We can now see, however, that for beliefs about the future expressed using bgt-sentences, this does not follow. Suppose Dodd believes that he is going to be in London. Given the linguistic fact just noted, this belief can be true even if Dodd is not in London in the future—it rather depends on how things stand right now. Here's an example.

Job Context

Dodd has been job hunting. Dave wants to know whether Dodd has found a job yet.

Dave: Dodd, where are you going to be next year? I myself am going to be in London. It would be good if we were working near to one another.

Dodd: Yes, I am going to be in London next year. Birkbeck have offered me a rather stable post.

Suppose Dodd dies before moving to London. Does this show that what he said is false? i.e. is it now false that he was going to be in London? No. It may well be false that he is now going to be in London—now that he's dead. But the claim that Dodd made is not shown false by his dying before next year.

Notice that for Dodd's utterance to be true even if he dies before next year, is not for Dodd's utterance to have the content that he'll be in London next year even if he dies before next year (i.e. dead or alive). This is important because, in Job Context, he obviously wasn't saying that he'll be in London next year dead or alive. So if the only way for his utterance to be true even if he dies before next year is for his utterance to have the content that he'll be in London next year dead or alive, then it cannot be correct to say—as we are—that Dodd's utterance could be true even if he dies before next year. Roughly speaking, Dodd's utterance is true if and only if something about the world at the time of utterance is such that, given that it is in place, Dodd is in London next year. This doesn't entail that he is in London next year because whatever it is about the present which makes it the case that next year he's in London, could be destroyed or otherwise cease to exist.

Here's an analogy. Suppose a ball is rolling down a chute and at the base of the chute is an upright domino. I say, "that ball is going to run into the domino." But suppose that at this moment, my friend intervenes and breaks the chute so that the ball doesn't run into the domino. Was what I said falsified? No. Why? Because the state of affairs in the present that made it true—i.e. made it the case that in the future the ball runs into the domino—was destroyed. But it nonetheless existed at the time I spoke. To falsify Dodd's utterance one would need to show that whatever it is about the present which he thinks makes it the case that next year he's in London doesn't in fact make this the case—so even if it did persist through to next year, it would still be open that Dodd is not in London next year. For instance, one would have to show that he hasn't actually been given the job that he thinks he has or that, as of some time prior to the time of utterance, Birkbeck relocated to somewhere outside of London, etc.

Let's turn now to Dodd's example of a belief about the future expressed using a *w*-sentence. We can bring out a feature of the semantic behaviour of *w*-sentences by once again contrasting *w*-sentences with corresponding *bgt*-sentences. Imagine each of the following two sentences written on a billboard advertisement as you drive along the highway, down which is located a town called Madera:

- (25) We'll change your oil in Madera.
- (26) We're going to change your oil in Madera. (Copley, 2004, 2009)

(26) sounds bullying whereas (25) sounds like an offer. (25) is conditional upon whether you, the reader, want your oil changed in Madera, whereas (26) is not—the decision has already been made to change your oil in Madera, whether you want it or not. The same contrast can be found in the following—suppose that each appears in a guide to gardening (adapted from (Palmer, 1979, p. 139)):

- (27) (a) Your nurseryman will spare you a few understocks.
(b) Your nurseryman is going to spare you a few understocks.
- (28) (a) Powdered chalk will act as a corrective of overly acidic soil.
(b) Powdered chalk is going to act as a corrective of overly acidic soil.
- (29) (a) Bridget will tell you that she was at the same lecture.
(b) Bridget is going to tell you that she was at the same lecture.

In each case, it seems that the (a) of each pair is implicitly conditional in a way that is made obvious when we contrast it with the (b) of the pair. When the gardening guide includes (27a), it is not telling you that something is definitely going to happen in the future—in the way that (27b) does. It's clear that it's telling you what is the case in the future where (amongst other things) you ask your nurseryman for understocks. Similarly, (28a) is not telling you what (28b) is telling you: viz. that in the future chalk is acting as a corrective to overly acidic soil. Rather, (28a) is telling you that chalk does this in the future where (amongst other things) you apply it in the correct way to the acidic soil. The same applies to (29a) and (29b). In general, *w*-sentences (formulated either with "will" or the

more archaic “shall”) are very frequently, implicitly conditional (Palmer, 1979, p. 138).

What a w-sentence is conditional upon depends upon the context of its use. Consider (30) uttered in two different contexts.

(30) You will die.⁷

Compare, (30) being said in a conversation about mortality versus (30) being said in a conversation about whether you are likely to survive a trip through the Colombian red zone. In the second context (30) may be false while in the first context it is true.

The sentence upon which Dodd focuses (a w-sentence) behaves similarly: it is implicitly conditional in a way that varies with context. Consider again:

(22) I shall not be able to afford an African safari.

It is easy enough to construct quite ordinary looking contexts in which (22) is used to say something true, even if Dodd wins the lottery.

New Job Context

Dodd is planning a holiday for next year. Dodd has a modest income. But he does regularly play the lottery. Dave knows this.

Dodd: Could you help me with something?

Dave: Sure.

Dodd: I'm trying to plan my holiday for next Summer. I'm not going to pay for it this year. I'll pay for it next year — after I get my new job. When I get my new job my pay goes up. So my budget for a holiday will then be X. So what do you think I will be able to afford on that budget?

Dave: Well, you will be able to afford a camping site in Bognor Regis.

Dodd: Hmm. But I shall not be able to afford an African Safari.

Dave: Well, that's not true. You might win the lottery.

Dodd's claim, made using (22), in this context, is not shown false by his winning the lottery; note the infelicity of Dave's last remark. Even if Dodd does win the lottery, that won't change the facts about what he can afford given his budget of X. All it does is change what his budget is. But given that here he is speaking about what he can afford given his then current budget, what he could afford given an alternative budget is irrelevant to the truth of the claim he makes here. Thus, just as one wouldn't show an utterance of (30) to be true when uttered in a context about the Colombian red zone by pointing out that everyone is mortal, so too, one wouldn't show an utterance of (22) to be not true, when

⁷ Notice that the context-sensitivity in sentences about future death allows for the possibility that one's evidence may well entail one's beliefs about one's future survival, and thus that Dodd's assumption that one's evidence won't entail that one is going to be, or will be, alive at some point in the future may be challenged on the same ground as that pursued currently.

uttered in the described context.

The following objection has been put to me: “well, having uttered (22), if *I* were asked, 'Do you think that what you said would be true even if you won the lottery?' then *I* would reply, 'No, then what I said would be false.' And this shows that what I was saying would be falsified if I won the lottery.” But I (the author) think this objection is raised in bad faith. If you openly admit that you play the lottery and that this could significantly change the funds available, and when you expressed your belief using (22), you weren't speaking with an appropriate restriction on the modal “shall”, then your belief doesn't even seem to constitute knowledge: so understood, there is no impression that you *know* that you shan't be able to afford an African safari. It's only if you have the restriction on the modal and, for that reason, deny that winning the lottery is relevant to the truth of what you were saying (just as mere mortality does not verify (30) when uttered in a conversation about the Colombian red zone), that your belief would, intuitively, seem to be knowledge.

Dodd therefore doesn't provide us with unequivocal examples of beliefs about the future which we take to be knowledge but which are not entailed by the evidence upon which they're based. The intuition that Dodd can know that he shall not be able to afford an African safari (even as he plays the lottery) and that Dodd can know various things about what he is going to be doing in the future (even though he is mortal) could easily be elicited by the fact that such beliefs are consistent with his winning the lottery and his premature death. For Dodd's examples to unequivocally support premise (1) of the example based objection, we need some assurance that such readings of sentences about the future are not the source of our intuition on this matter. Dodd provides no such assurance.

5 Providing assurance

I now outline two difficulties that will be faced by someone seeking to produce examples in support of premise (1) which clearly avoid the criticism just levelled against Dodd and Dutant.

The first difficulty is as follows. In sketching an example in support of premise (1) we can provide a context in which it is quite uncontroversial that the belief is not entailed by the evidence it is based upon. This will serve to block the infallibilist from being able to accuse the example of trading on a vacillation between different readings of the sentences in the example. However, it seems likely that the more the context makes it abundantly obvious that there is no entailment, the weaker will be the intuition that the belief constitutes knowledge.

There are, for example, some contexts in which the sentence “I am going to be in London next year” would be false if Dodd died prematurely. For instance:

Mortality Context

Dave: Tell me one thing that, regardless of whether you die this year, will be true

next year.

Dodd: Hm. Here's one: I am going to be in London next year.

Assuming that Dodd ceases to exist when he dies (or that he has no plans to be buried in London), Dodd's utterance is false if he dies before next year. I assume that the reader shares my judgement that if Dodd were to say the following in answer to Dave's question, he would say something that is not intuitively true:

Dodd: Hm. Well: I know that I'm going to be in London next year.

Similarly, here's a context in which Dodd's (22) is used to express a proposition that is false if Dodd wins the lottery:

Pub Context

Dave: Tell me something that you won't be able to afford next year, even if you win the lottery.

Dodd: I shall not be able to afford an African safari.

But in this context, if Dodd were to make a knowledge claim with this sentence, then the knowledge claim would not be intuitively true:

Dodd: I know that I shall not be able to afford an African safari.

It seems likely that attempts to make a non-entailing reading undeniably the right one for the context will generally suffer the same fate. Why think so? Concessive knowledge attributions are claims of the form "I know that p but it might be that q" (where q is obviously inconsistent with p). These are generally infelicitous. The contexts we're currently focusing upon are those in which it is abundantly obvious that a thinker believes a proposition which is obviously consistent with the thinker's evidence. It seems very likely that when this is very obvious, it'll have the same effect on a knowledge claim as the second conjunct of a concessive knowledge attribution. For instance, if you believe that it will rain tomorrow on the basis of a disreputable weather report website which is known to be highly unreliable—so it's obvious that your belief's falsity is consistent with your evidence—then the intuition that your belief constitutes knowledge instantly wanes. The more obvious we make the fact that your evidence is consistent with the falsity of what you believe, the more our judgements about knowledge will approximate what we see with concessive knowledge attributions. If this is right, then it seems likely that blocking the criticism I have been describing, by finding a suitable context, will be quite hard to accomplish.

The second difficulty is that context-sensitivity can show up in places where you don't expect it. For instance, very often, we justify a belief that p, with a sentence of the form "A told B that p." It is widely assumed that the latter sentence does not—on any reading—entail the truth of "p." So any filling-in of these schematic sentences will—one might

think—be an example that supports premise (1) of the example based objection against infallibilism. For instance, consider this: I believe that my mother lived in Switzerland when she was younger. I believe this because she told me that she lived in Switzerland when she was younger. Yet, surely, it's quite possible for her to tell me that she lived in Switzerland when she was younger even though she did not live in Switzerland when she was younger. But the widely held assumption is false.

As observed by Spector and Egré (2015, p. 1738), '*tell someone that S* has a *factive use*, though it is not *always* factive (as opposed, say, to *know*).'¹ Here's some evidence for this claim. Consider the following three sentences:

- (31) Sue told someone that she used to live in Switzerland.
- (32) Sue didn't tell anyone that she used to live in Switzerland.
- (33) Did Sue tell anyone that she used to live in Switzerland?

Each sentence has a reading on which it implies that Sue used to live in Switzerland. That implication thus exhibits a pattern typical of presupposition: i.e. it is preserved under negation and question-formation. This implication also passes the *Wait a minute!* test for presupposition described by von Stechow (2004). According to this test, if one can felicitously add to a sentence that purportedly presupposes that P, "Hey wait a minute, I didn't know that P" then that is reason to believe that P is presupposed by the relevant sentence. One can felicitously add, "Hey wait a minute! I didn't know that Sue used to live in Switzerland" to each of (31)-(33). So there is reason to accept Spector and Egré's claim. If we do then we should also grant that the proposition expressed by the sentence, "she told me that she lived in Switzerland when she was younger," as used in some contexts, will entail the proposition expressed by the sentence, "my mother lived in Switzerland when she was younger" as used in those contexts—contrary to first appearances. We should concede, that propositions expressed with the sentence "A told B that p" can indeed entail corresponding propositions expressed by the sentence "p."

Generally, if you think you have an example of a pair of sentences that can be used to express evidence and belief based upon that evidence such that the former doesn't entail the latter, then you need to ensure that the infallibilist cannot uncover context-sensitivity in these sentences that affords a reading of them on which there is an entailment. This requires more care than is typically given to the choice of examples given in support of premise (1) of the example based objection.

In sum, to produce adequate example-based support for premise (1) of the example based objection, two things need to be done. Firstly, all potentially relevant context-sensitivity in the sentences used to express evidence and evidence-based belief needs to be identified. Secondly the intuition that the belief constitutes knowledge must be elicited even while ensuring that no reading of the sentences is contextually available on which the evidence entails the belief. Although this may be possible to do, it is nonetheless significantly harder to accomplish than merely identifying some *sentences* that we use to express our evidence

and evidence-based beliefs.

6 A note on inductive belief

Although it has long been argued that perceptual beliefs can be entailed by the evidence that supports them (e.g. "I saw that p"), and although similar things can be said of testimony (e.g. "He told me that p" or "He let me know that p") and memory based beliefs (e.g. "I remember that p"), one might be tempted to think that inductive beliefs are never going to be saved for the infallibilist. Surely: a belief that is as general as an inductive belief but which is based upon particular observations, is never going to be shown to be entailed by the evidence upon which it is based.

However, this really depends upon the semantics of the sentences used to express inductive evidence-based beliefs and the evidence upon which such beliefs are based. Suppose I watch some water freeze when reduced to a temperature below 0 degrees centigrade. Maybe I see this happen *n* times and I form an inductive belief. Here are two ways to describe what happens. Firstly: what I observed was *n* samples of water freezing when their temperature was reduced below 0 degrees; what I believe is a universally quantified generalization over samples of water and circumstances: for all *x*, such that *x* is a sample of water, whenever the temperature of *x* is reduced below 0 degrees, *x* freezes. Alternatively: we can describe what I observed using a mass noun (and not a count noun phrase viz. "sample of water") and a verb with generic or habitual aspect (viz. "freezes" predicated of the mass noun): I observed that water freezes when its temperature is reduced below 0 degrees, or, I observed water freezing when its temperature is reduced below 0 degrees. What I believe on this basis is that water freezes when its temperature is reduced below 0 degrees. Given the first description of things, what I observe doesn't entail what I believe: since I have only observed *some* samples of water but my belief concerns *all* samples of water. But given the second description of things, what I observe could easily entail what I believe: I observe that water does something, and I believe that it does it.

Dodd was wrong to think that we can only report what one remembers using a past tense embedded clause (e.g. "I remember that a desk was in my office"). You can use the present tense too (e.g. "I remember that a desk is in my office"). This (given certain observations about the present tense) changes the prospects of what one can know about the world which one does not immediately observe, given infallibilism (at least for all Dodd has shown). Given the two descriptions of inductive belief just sketched, we should, in addition, acknowledge that there's the potential to say something similar in the case of inductive belief. It might well be that inductive belief can be known through a combination of observation and entailment—a possibility that has been explored to some extent by Everett (2010 esp. pp. 319-320).

This possibility is in fact implied by at least one respectable account of scientific law-statements and our use of them to think about the world they describe; namely, Nancy Cartwright's (1999) (see also (Hempel, 1988)). A traditional account of scientific law-

statements treats them as unbounded generalizations about the universe. We infer them from observations that do not entail them. We use them to make predictions about circumstances that may well be radically different from the observed circumstances upon which they are justified. But on Cartwright's account, law-statements are blueprints for types of complex object (which Cartwright calls "nomological machines"), which work in accordance with the relevant law-statement. We use a law-statement to reason about the world around us by learning how to identify or build a complex object whose internal components operate according to the law-statement. We then use our knowledge of the internal workings of the complex object (once we have identified it) to make predictions about the internal workings of that complex object. Such a process need involve no leap from observed evidence to a belief that is not entailed by the evidence. You observe that some water is part of a complex object such that if the water's temperature falls below 0 degrees, it will freeze. You observe that the water's temperature has fallen below 0 degrees. On the basis of these observations, you believe that the water will freeze. On such an account of law-statements and their use, apparently "inductive" belief can in principle be entailed by the evidence it is based on.

There is room, then, for the infallibilist to allow that we can have inductive knowledge, if the infallibilist is willing to explore less orthodox accounts of inductive belief and inductive reasoning. I won't pursue this possibility any further because, at least in this paper, I merely intend to draw attention to how the difficulties of providing examples in support of premise (1) of the example based objection are as likely to arise in the case of inductive belief as they are in the cases we have already considered. The truth-conditions of sentences used to express inductive belief (composed as they are of generics and verbs with habitual or generic aspect) are complex and quite probably, in some way, context-sensitive ((Carlson, 2012; Leslie, 2016)). Showing that there are no readings of such sentences that allow for an infallibilist story of inductive knowledge is not going to be simple.

References

- Bennett, M., & Partee, B. (1972). *Toward the Logic of Tense and Aspect in English*. Santa Monica, California: System Development Corporation.
- Bonjour, L. (2010). The Myth of Knowledge. *Philosophical Perspectives*, 24, 57–83.
- Carlson, G. (2012). Habitual and Generic Aspect. In R. Binnick (Ed.), *The Oxford Handbook of Tense and Aspect* (pp. 828–851). Oxford: Oxford University Press.
- Cartwright, N. (1999). *The Dappled World - A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.
- Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives*, 2, 91–123.
- Copley, B. (2004). Aspect, Scope, and Future Conditionals. *Proceedings of the Conference "Sub8 - Sinn Und Bedeutung,"* 177, 67–82.

- Copley, B. (2009). *The Semantics of the Future*. New York: Routledge.
- Davis, W. (2007). Knowledge claims and context: loose use. *Philosophical Studies*, 132(3), 395–438.
- Dodd, D. (2007). Why Williamson should be a sceptic. *The Philosophical Quarterly*, 57(229), 635–649.
- Dodd, D. (2011). Against Fallibilism. *Australasian Journal of Philosophy*, 89(4), 665–685.
- Dougherty, T. (2011). Fallibilism. In D. Pritchard & S. Bernecker (Eds.), *The Routledge Companion to Epistemology* (pp. 131–143). London: Routledge.
- Douven, I. (2005). Lewis on fallible knowledge. *Australasian Journal of Philosophy*, 83(4), 573–580.
- Dretske, F. (1970). Epistemic Operators. *Journal of Philosophy*, 67(24), 1007–1023.
- Dutant, J. (2007). The Case for Infallibilism. In C. Penco, M. Vignolo, & C. Amoretti (Eds.), *Proceedings of the 4th Latin Meeting in Analytic Philosophy* (pp. 59–84). Genoa: University of Genoa.
- Dutant, J. (2016). How to be an infallibilist. *Philosophical Issues*, 26(1), 148–171.
- Everett, T. (2010). Observation and Induction. *Logos and Episteme*, 1(2), 303–324.
- Fantl, J., & McGrath, M. (2009). Advice for fallibilists: put knowledge to work. *Philosophical Studies*, 142(1), 55–66.
- Fumerton, R. (2010). Fencing out Pragmatic Encroachment. *Philosophical Perspectives*, 24(1), 243–253.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Clarendon Press.
- Hempel, C. (1988). Provisos: a problem concerning the inferential function of scientific theories. *Erkenntnis*, 28, 147–164.
- Klein, W. (1994). *Time in Language*. London: Routledge.
- Leddington, J. (forthcoming). Fallibility for Infallibilists. In J. Gersel, R. Jensen, M. Thaning, & S. Overgaard (Eds.), *Experiential Reasons*. Oxford: Oxford University Press.
- Leslie, S. (2016). Generic Generalizations. In E. Zalta (Ed.), *Stanford Encyclopedia of Philosophy* (Winter 2016). Retrieved from <https://plato.stanford.edu/archives/win2016/entries/generics/>
- Lewis, D. (1996). Elusive Knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.

- Littlejohn, C. (2008). From $E = K$ to Scepticism? *The Philosophical Quarterly*, 58(233), 679–684.
- McDowell, J. (1982). Criteria, Defeasibility, and Knowledge. In *Meaning, Knowledge, and Reality* (pp. 369–394). Cambridge, MA: Harvard University Press.
- O'Brien, D. (2006). *An Introduction to The Theory of Knowledge*. Cambridge: Polity Press.
- Palmer, F. (1979). *Modality and the English Modals*. London: Longman.
- Partee, B. (1973). Some structural analogies between tenses and pronouns in English. *Journal of Philosophy*, 70(18), 601–609.
- Reed, B. (2012). Fallibilism. *Philosophy Compass*, 7(9), 585–596.
- Spector, B., & Egré, P. (2015). A uniform semantics for embedded interrogatives: an answer, not necessarily the answer. *Synthese*, 192(6), 1729–1784.
- Travis, C. (2005). A Sense of Occasion. In *Occasion-Sensitivity: Selected Essays* (pp. 290–315). Oxford: Oxford University Press.
- von Fintel, K. (2004). Would you Believe it? The King of France is Back! (Presuppositions and Truth-Value Intuitions). In M. Reimer & A. Bezuidenhout (Eds.), *Descriptions and Beyond* (pp. 315–341). Oxford: Clarendon Press.
- Wittgenstein, L. (2001). *Philosophical Investigations*. Oxford: Blackwell.