



Okasha, S. (2024). Scepticism, Evidential Holism and the Logic of Demonic Deception. *Noûs*. Advance online publication. <https://doi.org/10.1111/nous.12490>

Peer reviewed version

License (if available):
CC BY

Link to published version (if available):
[10.1111/nous.12490](https://doi.org/10.1111/nous.12490)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the accepted author manuscript (AAM) of the article which has been made Open Access under the University of Bristol's Scholarly Works Policy. The final published version (Version of Record) can be found on the publisher's website. The copyright of any third-party content, such as images, remains with the copyright holder.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available: <http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Scepticism, Evidential Holism and the Logic of Demonic Deception

I conjectured that physical theory, the global system of the world, is underdetermined, but not that every subordinate system was underdetermined.

W.V. Quine (1979)

Abstract

Sceptical arguments in epistemology typically employ *sceptical hypotheses*, which are rivals to our everyday beliefs so constructed that they fit exactly the evidence on which those beliefs are based. There are two ways of using a sceptical hypothesis to undermine an everyday belief, giving rise to two distinct sorts of sceptical argument: underdetermination-based and closure-based. However, both sorts of argument, as usually formulated in the literature, fall foul of evidential holism, for they ignore the crucial role of background beliefs. An analogy with the philosophy of science makes this point explicit. There is no simple way to “holism proof” the two sceptical arguments.

Keywords— scepticism, underdetermination, closure, justification, empirical equivalence, sceptical hypotheses

1 Introduction

A familiar pattern of sceptical argument in philosophy invokes underdetermination of our beliefs by our evidence. In such arguments, the sceptic tries to cast doubt on some of our everyday beliefs by devising an incompatible *sceptical hypothesis* which fits the evidence equally well. The challenge is then to rule out this sceptical hypothesis, or to explain why our everyday beliefs are preferable to it. Unless the challenge can be met, the sceptic argues, our beliefs lack justification and/or do not constitute knowledge.

The content of the sceptical hypothesis depends on which beliefs the sceptic wishes to undermine. Thus for beliefs about the external world, we have Descartes' hypothesis that an evil demon is causing my sensory experiences (or its modern incarnation, the brain-in-vat hypothesis). For beliefs about the past, we have Russell's hypothesis that the world came into existence five minutes ago complete with my apparent memory traces (Russell 1921). For beliefs about the future, we have the hypothesis that the regularities that have held up until now will suddenly break down. In each case, the hypothesis in question is incompatible with our everyday beliefs but fits the relevant evidence equally well, or so the sceptic claims.

Though sceptical arguments based on underdetermination are nothing new, recent epistemology has seen renewed interest in unpacking exactly how they work.¹ The main focus has been on external world scepticism, to which we shall confine our attention here, understood as the claim that our beliefs about the external world are not justified. Much has been made of a contrast between two (apparently) different routes to the sceptical conclusion, one that relies on underdetermination, the other on a closure principle. This is an important contrast, though it is misleadingly named since both routes rely on underdetermination considerations in a broad sense, and both make use of sceptical hypotheses. Despite this, we shall follow established terminol-

ogy in distinguishing between underdetermination-based and closure-based arguments for external world scepticism.

Epistemology is not the only branch of philosophy where the issue of underdetermination has arisen. In the philosophy of science, underdetermination has been widely discussed in relation to “empirically equivalent theories”, that is, pairs of scientific theories that are incompatible but make identical empirical predictions. According to the underdetermination argument for scientific anti-realism, empirically equivalent theories are equally supported by the evidence, hence epistemically on a par; so if we can assume that every scientific theory in principle has an empirically equivalent rival, it follows that we are unjustified in believing the truth of any scientific theory.

It is obvious that there is a close parallel between underdetermination arguments in epistemology and in philosophy of science. This is by no means a new point. However, two key morals regarding underdetermination from the philosophy of science literature, that came into focus in the 1990s, have not been taken account of in the epistemology literature. Both morals derive from evidential holism – the idea that our beliefs are responsible to empirical evidence not singly but in groups. The aim of this paper is to fill the lacuna and to trace the implications for the arguments for external world scepticism.

The structure of the paper is as follows. Section 2 lays out the contrast between underdetermination-based and closure-based arguments for scepticism about the external world. Section 3 discusses underdetermination arguments in the philosophy of science and the challenge posed to such arguments by evidential holism. Section 4 returns to epistemology and applies the lessons learned. Section 5 argues that there is no easy way to formulate “holism-proof” versions of the two sceptical arguments. Section 6 considers the epistemological implications of this finding.

2 Two arguments for external world scepticism

The underdetermination-based and closure-based arguments share a common conclusion, namely that our beliefs about the external world are unjustified. (Here I focus on justification-scepticism rather than knowledge-scepticism.²) Both arguments target, in the first instance, beliefs about one's immediate environment, for example, my belief that it is sunny outside. According to a venerable philosophical tradition, such beliefs are based on current sensory experiences, in the sense that it is those experiences (or on some versions, beliefs about the experiences), that constitute the evidence for the beliefs.³ Both arguments then introduce a sceptical hypothesis that is incompatible with the beliefs in question but is so chosen that the evidence cannot discriminate between them, for example that the evil demon is deceiving me into believing that it is sunny outside by causing my current sensory experiences. However the two arguments utilize the sceptical hypothesis differently to yield the sceptical conclusion.

Some notation will be useful. Let p denote one of my beliefs about the external world that the sceptic wishes to target. Let $\text{JB}(p)$ mean that my belief that p is justified. The sceptic's conclusion is thus $\neg\text{JB}(p)$. Let $D(p)$ denote the rival sceptical hypothesis; we may read this as "the demon is tricking me into believing that p by causing me to have sensory experiences qualitatively indistinguishable from the actual⁴ experiences on the basis of which I believe that p ". Note that D is a sentential operator.

Two points regarding the logic of the D operator will be important. Firstly, D is anti-factive, that is, $D(p) \Rightarrow \neg p$; for the demon can only *deceive* me into believing that p if p is false. Anti-factivity is essential to both sceptical arguments, as they require that the sceptical hypothesis $D(p)$ be incompatible with my belief that p . Secondly, D does *not* distribute across

conjunction. For suppose that it did, i.e. that $D(p \wedge q) \Rightarrow D(p) \wedge D(q)$. Given anti-factivity, it then follows that $D(p \wedge \neg q)$ is false whenever $q \Rightarrow p$; thus we can know *a priori* that certain sceptical hypotheses are false! For example, let p be “the ball is not green” and q be “the ball is red”. Presumably, it is possible that the demon could deceive me into believing $p \wedge \neg q$ (“the ball is neither green nor red”) by giving me the relevant sensory experiences, e.g. by making the ball appear yellow. But if D distributed across conjunction, I could immediately infer that this demon hypothesis is false.

2.1 The UN argument

The underdetermination-based argument, denoted **UN**, appeals to the subjective indistinguishability of p and $D(p)$. Since my actual sensory experiences, on the basis of which I believe that p , would be exactly the same if instead the rival hypothesis $D(p)$ were true, surely my evidence doesn’t favour p over $D(p)$? And if this is granted, it seemingly follows that I am not justified in believing that p , given that p and $D(p)$ are incompatible. The **UN** argument can thus be expressed as follows:

UN1: if a person’s evidence doesn’t favour x over an incompatible y , then $\neg\text{JB}(x)$

UN2: my evidence doesn’t favour p over $D(p)$

therefore

C: $\neg\text{JB}(p)$

Note that **UN** is strictly speaking an argument-schema, instances of which

are formed by substituting a suitable proposition for p in **UN2**.

Premise **UN1** is to be read as quantifying over all persons and propositions; it says that if propositions x and y are incompatible, a person is not justified in believing x unless their evidence favours x over y . Plausibly, this expresses a conceptual truth about the relation between evidence and epistemic justification, at least if we accept a broadly “evidentialist” view of justification. **UN1** is a standard formulation of what a number of authors call the “underdetermination principle”, though this term will not be used here.⁵

I will assume that **UN1** is acceptable. But one issue merits brief comment, which is what “favouring” means. McCain (2013) writes, in the context of the **UN** argument: “for one’s evidence to favor x over y is simply for one’s evidence to support x to a higher degree than it does y ” (p. 291n).⁶ This sounds right; but there is an ambiguity here, that is made manifest when we express the concept of support in probabilistic terms. If evidence e supports x to a higher degree than it does y , does this mean that $Pr(x | e) > Pr(y | e)$, or does it mean that $Pr(e | x) > Pr(e | y)$? The latter corresponds to the “likelihoodist” conception of support defended by certain philosophers and statisticians.⁷ However the former is the reading required to vindicate **UN1**. It is very plausible that if x and y are incompatible, I am justified in believing x only if $Pr(x | e) > Pr(y | e)$, where e is my total evidence; but it is not at all plausible that I am justified in believing x only if $Pr(e | x) > Pr(e | y)$.

What about premise **UN2**, the claim that my evidence doesn’t favour p over $D(p)$? This premise depends on a background epistemological assumption and a principle about evidential support. The background assumption is the one mentioned previously, namely that sensory experiences constitute the evidence for our beliefs about the external world. The principle is that where two rival hypotheses (or beliefs) both fit all the evidence, they are equally supported by that evidence, or evidentially on a par; and hence the evidence

doesn't favour one over the other. Let us call this the *Parity* principle.

What exactly does “fit” mean here? In one straightforward sense, what it is for a hypothesis to fit the evidence is for it to imply that evidence. But a potential objection to **UN2** then looms, which is that p and $D(p)$ do *not* both fit the evidence. For although $D(p)$ implies something about my current sensory experiences, p does not. That the demon is tricking me into believing that it is sunny outside by causing me to have certain sensory experiences implies that I do in fact have those experiences; but that it is sunny outside implies nothing about what sensory experiences, if any, I currently have. This prompts the worry that the way the **UN** argument is often set up in the literature, in terms of an opposition between p and $D(p)$, is not quite right. We shall return to this worry below.

2.2 The CL argument

The closure-based argument, denoted **CL**, proceeds differently. In this argument, the sceptic claims that I am not justified in believing that the sceptical hypothesis $D(p)$ is false. She then invokes the principle that justified belief is closed under known logical consequence, and the fact that p and $D(p)$ are known to be incompatible, to yield the conclusion that I am not justified in believing that p . The **CL** argument can thus be expressed as follows:

CL1: if $\text{JB}(x)$, and it is known that $x \Rightarrow y$, then $\text{JB}(y)$

CL2: $\neg\text{JB}\neg D(p)$

therefore

C: $\neg\text{JB}(p)$

As before, **CL1** is understood as quantifying over all persons and propositions; and **CL2** is to be filled with appropriate choice of p .

CL1 is a closure principle for justified belief. It is widely thought that some such principle must be right, though the optimal formulation is a matter of debate; this issue will not matter here. **CL2** says that I am not justified in believing the falsity of the sceptical hypothesis $D(p)$. The motivation for **CL2** is that my evidence doesn't rule out $D(p)$, since if $D(p)$ were true then it would explain my current sensory experiences; thus I am unjustified in believing that $D(p)$ is false. Cohen (1998) suggests that the operative principle here is something like: if e is my total evidence for x , and there is an x -incompatible y which if true would explain e , then I am not justified in believing $\neg y$ (p.147n). Let us call this the *Caro* principle (for "Can't rule it out").⁸

Note that the *Caro* principle only establishes **CL2** if we assume that my current sensory experiences constitute the total evidence for my belief that p . For otherwise, the fact that $D(p)$, if true, would explain those experiences does not show that I am unjustified in believing $\neg D(p)$, even granting *Caro*. Thus the **CL** argument, like the **UN** argument, rests on the epistemological assumption that the relation of sensory experience to worldly belief is that of evidence to hypothesis.

If **CL2** is indeed motivated by the *Caro* principle, an objection looms.⁹ For note that the sceptical hypothesis $D(p)$ is in effect a conjunction: "the demon is causing my sensory experiences..." and "not- p ". It is the first conjunct that explains my experiences, but the second conjunct that is incompatible with p . Thus in so far as explaining my experiences is concerned, $D(p)$ contains an irrelevant conjunct. And we know that adding irrelevant conjuncts tends to destroy explanation (Salmon 1984). So it is not true that $D(p)$ itself explains my experiences; rather a proper part of it does, and that

part is perfectly compatible with p . So the *Caro* principle does not imply **CL2** after all.

We shall return to this objection in section 5. For now, note that the objection does not work for all choices of p . If p is “it is sunny outside”, then the first conjunct of $D(p)$ does indeed explain my experiences and is compatible with p . But if p is “I can see the sun”, this is not so. For to see the sun requires that the sun be the cause of my sensory experiences and thus that the demon not be their cause. So for this choice of p , the sceptical hypothesis $D(p)$ does not factorize into two components where the first does the explanatory work and yet is compatible with p .

2.3 Discussion

What is the logical relation between the **UN** and **CL** arguments? A number of authors have argued that **CL1** implies **UN1**, given certain plausible principles about justified belief; but the converse inference is a matter of dispute.¹⁰ But whatever one’s view on this matter, the two arguments still offer different routes to the sceptical conclusion (see Figure 1). For the second premises of each argument, **UN2** and **CL2**, appear to be logically independent of each other and also of both of the first premises. Thus irrespective of the relation between **UN1** and **CL1**, the two sceptical arguments are still distinct. This at least will be our working assumption here.

Many epistemologists regard the **UN** and **CL** arguments as “recipes” for generating a sceptical conclusion about virtually any of our worldly beliefs, thus giving rise to a far-reaching form of scepticism. For example, Pritchard (2005) says that both arguments “can be repeated with just about any everyday proposition (one would just have to vary the sceptical hypothesis to suit)....and any subject”, and hence that “radical scepticism results” (p.40).

That the **UN** and **CL** arguments generalize in this way *seems* plausible.




Figure 1: **Two sceptical arguments**

For although the arguments target, in the first instance, beliefs that are based on current sensory experience, they could presumably be adapted to beliefs based on other sources, for example past sensory experiences, induction, or memory. For the evil demon could equally have caused my past sensory experiences, or have planted in me (apparent) memory traces of events that did not occur. So if either form of sceptical argument succeeds in undermining my belief that it is sunny outside, a parallel argument will presumably apply to all of my beliefs about the external world. We return to this issue later.

There are various ways of trying to resist the two sceptical arguments. One might reject the evidentialist conception of justification that underpins **UN1**, either by adopting reliabilism or some form of “immediate justification”.¹¹ One might reject the *Parity* principle that underpins **UN2**, for example by arguing that my belief that p constitutes a “better explanation” of my sensory experience than does the sceptical hypothesis $D(p)$ (perhaps because it is simpler), so my evidence does favour p over $D(p)$.¹² One might reject the closure principle for justified belief enshrined in **CL1**, as some philosophers have done.¹³ One might query the *Caro* principle that underpins **CL2**, or dispute **CL2** directly, as for example Pryor (2000) does. Finally,

one might try to dispose of both the **UN** and **CL** arguments in one swoop by rejecting the underlying epistemological assumption that sensory experience constitutes the evidential basis for beliefs about the external world.¹⁴

Whether any of these anti-sceptical moves works is not my concern here, so for the sake of argument I set them aside. I will assume that **UN1** and **CL1** are acceptable; I will grant for the sake of argument that sensory experience constitutes the evidential basis for our worldly beliefs, and I will set aside the “best explanation” response. (Though explanatory considerations will be indirectly relevant, given that the sceptic’s case for **CL2** relies on the *Caro* principle.) My point is that even granting all this, both sceptical arguments, as standardly formulated, face a serious objection. Whether it is possible to re-formulate the arguments to avoid the objection requires careful investigation. To see the objection, a brief detour into the philosophy of science is needed.

3 Underdetermination in the philosophy of science

The underdetermination argument for anti-realism says that for any scientific theory, there is an incompatible empirically equivalent rival, i.e. one that has exactly the same empirical implications, hence fits the same evidence.¹⁵ (An “empirical implication” is a statement whose truth-value can be determined directly from observation or experiment.) Next, it is argued that empirically equivalent rivals are equally supported by the evidence; from which it is concluded that we are unjustified in believing that any theory is true. This argument relies on the *Parity* principle (empirical equivalence implies evidential parity), and is structurally similar to the **UN** argument for external world scepticism.¹⁶

Two well-known responses are firstly, to deny that scientific theories *do* face empirically equivalent rivals (or to argue that this is only true in a trivial sense); and secondly, to argue that empirically equivalent rivals are not necessarily equally supported, since they may differ in respect of “extra-empirical virtues” such as simplicity. Here I set aside these responses to focus on a different objection to the underdetermination argument, namely that it stands in tension with the doctrine of evidential holism.

According to evidential holism, our beliefs about the world face the “tribunal of experience”, as Quine (1951) put it, not singly but in relatively large blocks. That is, it is sets of propositions, not single propositions, that are confirmed or disconfirmed by empirical evidence; and the empirical support cannot generally be parcelled out among them. Therefore, the empirical justification for any single worldly belief typically derives from the empirical justification enjoyed by a larger set of beliefs to which it belongs.

Holism is at its most compelling when applied to highly theoretical beliefs, for example that electrons are negatively charged. It seems clear that the empirical justification for this belief stems from its belonging to a broader body of physical theory that enjoys empirical support. Certainly, we can point to particular experimental findings that support the belief that electrons are negatively charged, such as Millikan’s oil-drop experiments of 1911, but they do so against a backdrop of other theoretical beliefs, e.g. about x-rays, ionization, and the workings of Millikan’s apparatus. Whether holism applies to all beliefs is perhaps less clear. Quine himself argued that it is a matter of degree: some beliefs are located to the interior of our “web” while others are closer to the experiential boundary, with fewer inferential connections to others. This is plausibly true; for there seems to be no principled way of identifying a class of beliefs, scientific or quotidian, which are such that the empirical data that support those beliefs do so entirely independently of background beliefs.

Here I assume that evidential holism, as characterized above, contains a broad measure of truth. (This is not to endorse the more extreme versions of holism that have sometimes been touted, e.g. that it is always possible to hold onto any belief in the face of any empirical findings, or that one cannot test a single hypothesis without testing the whole of science.)

Given holism, we can still talk about the evidence for a single belief, in a derived sense. Suppose x is a member of a larger set of beliefs $\{x, y, z.. \}$ which has evidential basis $e = \{e_1, e_2, e_3 \dots \}$. Holism says that we cannot identify any single e_i as the evidence for belief x . However, we may take the total evidence for x to be the whole set e . That is, the total evidence for a single belief is the evidence that supports a larger cluster to which it belongs.

Holism has two (related) morals for the underdetermination argument, as philosophers of science have long recognized.¹⁷ The first moral is that the relata of “is empirically equivalent to” cannot be single scientific hypotheses but must be larger clusters. For empirical equivalence means “having the same empirical implications as”; and holism teaches us that single hypotheses typically do not have empirical implications of their own, unless conjoined with further assumptions. For example, Newton’s laws of mechanics and gravitation imply nothing about the planetary orbits that will be observed, until conjoined with the auxiliary assumption that non-gravitational forces are negligible. So non-trivial instances of empirical equivalence will have to involve clusters of scientific hypotheses sufficiently inclusive to have empirical implications of their own.

The second moral is this. Even if two scientific theories t_1 and t_2 are empirically equivalent, it is quite possible that when they are conjoined with other scientific propositions (perhaps yet to be discovered), the equivalence will be broken. That is, there may exist a proposition x , which there is independent reason to believe, such that $(t_1 \wedge x)$ and $(t_2 \wedge x)$ have different (possibly incompatible) empirical implications. This reflects an obvious

logical point, namely that the conjunction of two propositions may have empirical implications that are not implications of either proposition alone. This in turn shows that the inference from empirical equivalence to evidential parity can fail. Even though t_1 and t_2 are empirically equivalent, our (total) evidence may still discriminate between them, for example if $(t_1 \wedge x)$, but not $(t_2 \wedge x)$, makes an empirical prediction that turns out to be true. That is, we could have indirect empirical reasons to favour t_1 over t_2 , or vice-versa. Only if there is reason to think that no such x will be found, is the inference from empirical equivalence to evidential parity watertight.

Taken together, these two holistic morals mean that the underdetermination argument for scientific anti-realism is less straightforward than it seems. The first moral shows that to run the underdetermination argument, it is essential to focus on a reasonably inclusive body of theory; the argument cannot be run on a hypothesis-by-hypothesis basis. The second moral shows that even then, the inference from empirical equivalence to evidential parity can be queried (and without the problematic appeal to “extra-empirical virtues”), given that scientific knowledge forms an integrated totality. It is for essentially this reason that Quine (1979) restricted his underdetermination thesis to “physical theory, the global system of the world”, while conceding that “not every subordinate system [is] underdetermined” (p.66).

4 Application to scepticism

Let us return to the two arguments for external world scepticism. Consider the **UN** argument first. Premise **UN2** states that my evidence doesn’t favour my belief that p over the rival sceptical hypothesis $D(p)$ from which it is subjectively indistinguishable. This closely parallels the scientific anti-realist’s claim that the evidence does not favour one theory over an empirically equivalent rival; where empirical equivalence corresponds to subjective indistin-

guishability. The only difference is that in the scepticism case, the evidence in question is a single individual’s sensory experience, while in the philosophy of science case, it is publicly-accessible empirical data. Aside from this difference, the arguments are alike.

However, the sceptic’s case for **UN2** conflicts with the holistic lessons above. The sceptic assumes that for each of my worldly beliefs to which the sceptical argument applies – the possible substitution-instances for p in **UN2** – there exists a proprietary set of sensory experiences which constitutes its evidential basis, independently of other beliefs. But holism teaches us that this is not generally so.

To flesh this out, consider the two holistic morals in turn. The first moral was that the relata of “is empirically equivalent to” cannot usually be individual hypotheses. We have actually already seen a glimmer of this moral as it applies to the **UN** argument, in our observation above that $D(p)$ does, but p does not, imply that I have the current sensory experiences that I actually have. This means that strictly speaking, p and $D(p)$ cannot be treated as empirically equivalent rivals with respect to the data of sensory experience; which is how the sceptic tries to treat them. That is, p and $D(p)$ do not coincide in their “experiential implications”, for p has none.

Now there is an obvious fix for the sceptic, which is to replace the target belief p with a slightly different one. Define p^+ as: “ p , and the fact that p is the cause of my actual sensory experiences”. (Both Yalçın (1992) and Vogel (2004) take the target of the **UN** argument to be causal beliefs of this sort.) This yields a well-posed underdetermination problem, since p^+ and $D(p)$ are now alternative explanations of the same sensory experiences, and do coincide in their experiential implications. The conclusion of the **UN** argument is then $\neg JBp^+$. Although this does not imply the original sceptical conclusion $\neg JBp$, it is nonetheless a worrying conclusion in its own right, since p^+ is presumably something that I do believe.¹⁸

To appreciate why the **UN** argument needs to take p^+ rather than p as its target, consider a simple Bayesian analysis. Let e be a proposition describing my current sensory experience. Then, both p^+ and $D(p)$ imply e . Therefore, if e favours p^+ over $D(p)$, i.e. if $Pr(p^+ | e) > Pr(D(p) | e)$, this can only be because p^+ has a higher prior probability than $D(p)$ (by Bayes' theorem). Absent some reason to assign p^+ a higher prior than $D(p)$, it follows that the evidence e does not favour p^+ over $D(p)$ – just as premise **UN2** says. But this reasoning does not go through for p , since p does not imply e to begin with.

In short, by replacing p with p^+ in the **UN** argument, the sceptic can respect the first holistic moral, albeit at the price of a slightly different sceptical conclusion. But the second moral is a different story.

The second moral, to recall, was that there could be indirect empirical reasons to favour one hypothesis over an empirically equivalent rival, if the former but not the latter is integrated into a broader system of beliefs that enjoys empirical support. This point applies to everyday beliefs no less than scientific beliefs. Consider again my belief p that it is sunny outside. Though p is based fairly directly (we may assume) on my current sensory experience, it is also part of a broader system of beliefs. (Similarly for p^+) These include: that viewing conditions are normal; that under normal conditions how things appear is a reliable guide to how they are; that the sun is the sole source of illumination in my room; that I am in room; that there is a sun; that I inhabit a world of three-dimensional physical objects; and more. It is in the context of these background beliefs that my sensory experiences lead me to believe that p . In so far as these background beliefs are themselves well-supported by evidence, this means that my total evidence does in fact favour p over the sceptical hypothesis $D(p)$. For $D(p)$ does not similarly mesh with my background beliefs, and indeed may well be inconsistent with them.¹⁹

It follows that the sceptic's argument for **UN2** fails (even if we replace p

with p^+). The sceptic's strategy – trying to undermine my actual belief by constructing a rival hypothesis $D(p)$ that fits my sensory experiences equally well – would only work if there was a one-one evidential relation between beliefs and episodes of sensory experience. But holism teaches us that there is not. So the sceptic has not shown that p and $D(p)$ are equally supported by the evidence (nor p^+ and $D(p)$).

Now it might be replied that this holistic point, even if correct, is of little dialectical significance in the context of debating scepticism. For my background beliefs are themselves about the external world, so ultimately (we may assume) derive their justification from (other episodes of) sensory experience. (Recall that we are granting the sceptic their assumption that sensory experience constitutes the evidence for worldly beliefs.) So even if the holist is right that the true evidential situation is not “present experiences support p ”, but rather “past and present experiences support p plus background beliefs”, this may seem immaterial. For the sceptic can simply extend the scope of their argument to cover background beliefs too.

This reply is tempting, but care is needed. We need to consider *exactly* how the external world sceptic can establish their conclusion while respecting evidential holism. It is essential to avoid hand-waving at this juncture if we are to attain a proper understanding of how, if at all, the **UN** argument can be rendered holism-proof. That is the task of the next section. But firstly let us consider how holism bears on the closure-based sceptical argument **CL**.

Since the **CL** argument does not rely on empirical equivalence, there is no need to change the target from p to p^+ (though we can do if we wish). But the second holistic moral undermines the sceptic's case for premise **CL2** – which states that I am not justified in believing $\neg D(p)$. Recall that **CL2** is underpinned by the *Caro* principle: if e is my total evidence for x , and there is an x -incompatible y which if true would explain e , then I am not justified in believing $\neg y$. But holism means that my current sensory experiences do

not exhaust the total evidence for my belief that p . For p is integrated into a larger cluster of beliefs, so the total evidence for p comprises all those sensory experiences that constitute the evidence for the cluster. Therefore the fact that $D(p)$, if true, would explain my current sensory experiences does not imply that I am not justified in believing $\neg D(p)$, even granting *Caro*.

To conclude: the **CL** argument, no less than the **UN** argument, presupposes a one-one relation between episodes of sensory experience and individual beliefs, and thus flies in the face of evidential holism.

5 Can we holism-proof the sceptical arguments?

To holism-proof the **UN** and **CL** arguments, we need somehow to take account of background beliefs. The obvious solution is to apply the demon operator not just to the target belief but to the background beliefs too.

As before, let p be the belief that the sceptic wishes to target. Let $bb \equiv b_1 \wedge \dots \wedge b_n$ denote the conjunction of all those background beliefs that are indirectly evidentially relevant to p , in the sense described above. Given the epistemological assumption on which the sceptic relies, the evidential basis for bb consists of past sensory experiences. The exact content of bb does not matter greatly, though we will assume that bb includes the proposition that under normal conditions, external objects will cause corresponding sensory experiences, e.g. looking outside on a sunny day will cause sun-like sensory experiences. Thus we may take it that bb implies the occurrence of those past sensory experiences that form its evidential basis; and the conjunction $p \wedge bb$ implies the occurrence of my current sensory experiences. (In this way we can accommodate the first holistic moral without having to modify the target belief to p^+ , since $p \wedge bb \Rightarrow p^+$).

Now in principle, the demon could induce my past sensory experiences no less than my current sensory experiences, as noted previously. So we can

apply the demon operator to bb , where $D(bb)$ means that the demon deceives me into believing bb by causing the (actual) past sensory experiences on which bb is based.

There are then two possible ways for the sceptic to proceed, depending on whether the demon operator is applied before or after bb is conjoined with p . There is also a third, hybrid way. Let us examine them in turn.

5.1 Attempt 1: demonize-then-conjoin

Suppose that the sceptic applies the demon operator to both p and bb to yield a pair of sceptical hypotheses: $D(p)$ and $D(bb)$. (“The demon is deceiving me into believing p , and the demon is deceiving me into believing bb ”). The sceptic then conjoins these hypotheses to yield a single complex sceptical hypothesis: $D(p) \wedge D(bb)$. This can be thought of as a rival to my actual belief $p \wedge bb$. (We assume that I do actually hold the belief that $p \wedge bb$.)²⁰

Next, the sceptic runs the **UN** argument, but with $p \wedge bb$ substituted for p , and $D(p) \wedge D(bb)$ substituted for $D(p)$, in premise **UN2**. This gives:

UN1: if a person’s evidence doesn’t favour x over an incompatible y , then $\neg\text{JB}(x)$

UN2*: my evidence doesn’t favour $p \wedge bb$ over $D(p) \wedge D(bb)$

therefore

C*: $\neg\text{JB}(p \wedge bb)$

What should we make of this argument? The first thing to note is that

the conclusion \mathbf{C}^* says that I am not justified in believing $p \wedge bb$, while the conclusion \mathbf{C} of the original sceptical argument was that I am not justified in believing p . Now \mathbf{C}^* is less troubling a conclusion than \mathbf{C} , as we will see later. But this issue may be sidelined for now, for the sceptic *appears* to have a holism-proof route to the original sceptical conclusion via the closure-based argument \mathbf{CL} .

Starting from \mathbf{CL} , then substituting $p \wedge bb$ for p and $D(p) \wedge D(bb)$ for $D(p)$ in $\mathbf{CL2}$, gives the following:

$\mathbf{CL1}$: if $\text{JB}(x)$, and it is known that $x \Rightarrow y$, then $\text{JB}(y)$

$\mathbf{CL2}^*$: $\neg \text{JB} \neg (D(p) \wedge D(bb))$

therefore

\mathbf{C}^* : $\neg \text{JB}(p \wedge bb)$

Premise $\mathbf{CL2}^*$ says that I am not justified in believing that the complex sceptical hypothesis $D(p) \wedge D(bb)$ is false, while the conclusion \mathbf{C}^* says that I am not justified in believing $p \wedge bb$. But notice that instead of moving from $\mathbf{CL1}$ and $\mathbf{CL2}^*$ to \mathbf{C}^* , the sceptic can instead move directly to the original sceptical conclusion \mathbf{C} , namely $\neg \text{JB}(p)$. For note that $p \Rightarrow \neg D(p) \Rightarrow \neg D(p) \vee \neg D(bb)$. Since by de Morgan's law, $\mathbf{CL2}^*$ is equivalent to $\neg \text{JB}(\neg D(p) \vee \neg D(bb))$, it follows that $\mathbf{CL1}$ and $\mathbf{CL2}^*$ imply $\neg \text{JB}(p)$. So it seems as if the closure-based argument, at least, may be rendered holism-proof.

However, there is a catch. For the sceptic is *not* entitled to premise $\mathbf{CL2}^*$ of the closure-based argument. Recall how the argument for this premise is supposed to go: since the sceptical hypothesis $D(p) \wedge D(bb)$ explains the total evidence for my belief p , I am not justified in believing the falsity of that

hypothesis – since I “can’t rule it out”. But it is not necessarily true that $D(p) \wedge D(bb)$ explains the total evidence for p . (Recall that the total evidence for p is, by definition, the total evidence for $p \wedge bb$). Certainly, $D(p)$ explains the sensory experiences on the basis of which I believe p , and $D(bb)$ explains the past sensory experiences on the basis of which I believe bb ; but those past and present sensory experiences may not exhaust the total evidence for $p \wedge bb$, precisely because of holism.

An example will make this clear. Let p be the proposition that it’s now raining, and b_1 the background belief that whenever it rains there’s a rainbow visible in the sky. Suppose that I believe p on the basis of my rain-like sensory experiences at the present time; and I believe b_1 on the basis of (induction over) my rain-experiences and rainbow-experiences at earlier times. (We may suppose that I live in a particularly rainbow-prone region of the world, so my past rain-experiences and rainbow-experiences have always co-occurred.) Now clearly, $D(p)$ explains the experiences on which p is based, and $D(b_1)$ explains the experiences on which b_1 is based. But consider my conjunctive belief $p \wedge b_1$. Notice that $p \wedge b_1$ implies that there is a rainbow visible in the sky now, and thus that if I now turn my gaze skyward, I will have a rainbow-experience. The sceptical hypothesis $D(p) \wedge D(b_1)$ does *not* have this implication. So if I do turn my gaze skywards, and do have a current rainbow experience, this experience forms part of my total evidence for $p \wedge b_1$, and thus for p , but it is not explained by $D(p) \wedge D(b_1)$. Therefore, the sceptic’s case for premise **CL2*** fails.

Similar reasoning shows that the sceptic’s case for **UN2*** fails too. Recall the supposed argument for **UN2***: my actual belief $p \wedge bb$ and the complex sceptical hypothesis $D(p) \wedge D(bb)$ are empirically equivalent, hence evidentially on a par. But the claim of empirical equivalence is untenable. The conjunction of two beliefs may have experiential implications that go beyond what either of them has alone; so the experiential implications of $D(p) \wedge D(bb)$

may be a proper subset of those of $p \wedge bb$, as in the example above.

The first way of trying to holism-proof the two sceptical arguments therefore fails.

5.2 Attempt 2: conjoin-then-demonize

Suppose instead that the sceptic applies the demon operator to the conjunction $p \wedge bb$, rather than to p and bb separately, to yield a single sceptical hypothesis $D(p \wedge bb)$ – “the demon is deceiving me into believing $p \wedge bb$.” Then, the claim of empirical equivalence between $p \wedge bb$ and $D(p \wedge bb)$ does go through, so the sceptic can truly claim that these are rival hypotheses that explain the very same sensory experiences.

Consider then the **UN** argument, with $p \wedge bb$ substituted for p and $D(p \wedge bb)$ for $D(p)$:

UN1: if a person’s evidence doesn’t favour x over an incompatible y , then $\neg\text{JB}(x)$

UN2**: my evidence doesn’t favour $(p \wedge bb)$ over $D(p \wedge bb)$

therefore

C*: $\neg\text{JB}(p \wedge bb)$

This argument is an improvement, since the sceptic’s case for **UN2**** does not now fall foul of holism. However, the conclusion **C*** is not especially troubling. For bb is itself a long conjunction, so $\neg\text{JB}(p \wedge bb)$ says that I am not justified in believing a long conjunction of propositions, of which one

conjunct is p and the other conjuncts are my background beliefs $b_1 \dots b_n$. But arguably, this is compatible with each of my individual beliefs being justified. For as the lottery paradox shows, it is not obvious that justified belief obeys the conjunction principle: $\text{JB}(x) \wedge \text{JB}(y) \Rightarrow \text{JB}(x \wedge y)$. If we do not assume this principle, then the sceptic's conclusion that $\neg \text{JB}(p \wedge bb)$ is compatible with my justifiably believing p and each of $b_1 \dots b_n$. Moreover, even if we do accept the conjunction principle for justified belief (so favour a different route out of the lottery paradox), then from $\neg \text{JB}(p \wedge bb)$ it only follows that at least one of my beliefs $\{p, b_1, \dots, b_n\}$ is unjustified, but not that any particular belief is unjustified. So it does not follow that my belief that p is unjustified.

In short, \mathbf{C}^* is logically weaker than, and considerably less threatening than, the conclusion \mathbf{C} of the original sceptical argument. Moreover, if \mathbf{C}^* is the most that the sceptic can establish while respecting holism, the resulting dialectical situation is rather different from the one depicted in the literature. For as noted, epistemologists typically regard both the \mathbf{UN} and \mathbf{CL} arguments as recipes for generating a sceptical conclusion about *any* of our external world beliefs; but this is not so, if \mathbf{C}^* is the best the sceptic can do.

A related consideration is this. Plausibly, for many of my “everyday” beliefs about my immediate environment, the relevant set of background beliefs, *modulo* which my sensory experience supports those beliefs, will share a common member, namely that how things appear is a reliable guide to how they are; call this belief b_r . Then, when the sceptic applies the \mathbf{UN} argument to each everyday-belief-plus-background-beliefs, this will generate a series of conclusions of the form \mathbf{C}^* . Now even if we *do* accept the conjunction principle for justified belief, one way that these sceptical conclusions could be jointly true is if the belief b_r is unjustified but each everyday belief is justified. Again, this is substantially different from what the sceptic's argument is

usually taken to show, namely that none of our worldly beliefs is justified.

Can the sceptic do better by applying the conjoin-then-demonize strategy to the closure-based argument **CL**? Substituting $p \wedge bb$ for p and $D(p \wedge bb)$ for $D(p)$ in the second premise of **CL** gives:

CL1: if $\text{JB}(x)$, and it is known that $x \Rightarrow y$, then $\text{JB}(y)$

CL2**: $\neg\text{JB}\neg(D(p \wedge bb))$

therefore

C*: $\neg\text{JB}(p \wedge bb)$

Note that the case for **CL2**** does not fall foul of holism, unlike the case for **CL2***, since the sceptical hypothesis $D(p \wedge bb)$, if true, would indeed explain the total evidence for $(p \wedge bb)$. However, **C*** is the same relatively unthreatening conclusion discussed above. But might there in addition be a route from **CL1** and **CL2**** to the original sceptical conclusion **C**, i.e. to $\neg\text{JB}(p)$? Recall from Attempt 1 that there was a route from **CL1** and **CL2*** to $\neg\text{JB}(p)$, based on the fact that $p \Rightarrow \neg(D(p) \wedge D(bb))$. However, that route is closed off once **CL2*** is replaced with **CL2****, since p does *not* imply $\neg(D(p \wedge bb))$. In short, if the sceptic uses conjoin-then-demonize, rather than demonize-then-conjoin, then this serves to holism-proof the second premise of the closure-based sceptical argument, but it also means that the argument will only yield the conclusion **C***, rather than in addition yielding **C**.

Note that there would be a route from **CL1** and **CL2**** to **C** if the demon operator distributed across conjunction (or were closed under logical consequence). For then we could argue as follows: $D(p \wedge bb) \Rightarrow D(p)$; so by contraposition, $\neg D(p) \Rightarrow \neg D(p \wedge bb)$; so by **CL1**, $\text{JB}(\neg D(p)) \Rightarrow \text{JB}(\neg(D(p \wedge$

bb); so by contraposition and **CL1**, $\neg\text{JB}(\neg D(p \wedge bb)) \Rightarrow \neg\text{JB}(\neg D(p)) \Rightarrow \neg\text{JB}(p)$; therefore **CL1** and **CL2**** imply **C**.

However, as emphasized in section 2, the demon operator does not distribute across conjunction, so it is not true that $D(p \wedge bb) \Rightarrow D(p)$. One way to see this is to note that $D(p \wedge bb)$ is compatible with the truth of p but $D(p)$ is not. For by anti-factivity, $D(p \wedge bb) \Rightarrow \neg(p \wedge bb)$, while $D(p) \Rightarrow \neg p$. That is, if the demon is deceiving me into believing that $p \wedge bb$, it follows that $p \wedge bb$ is false, but not that p itself is false. So there exists a possible situation in which $D(p \wedge bb)$ and p are both true, but in which $D(p)$ is therefore false.

Might the sceptic concede the point that the demon operator cannot generally distribute across conjunction, but nonetheless argue that in many cases, a circumstance in which $D(p \wedge bb)$ is true will be one in which $D(p)$ is true too? A relevant point here is that the hardest-to-refute form of external world scepticism arguably stems from “radical” sceptical scenarios like the brain-in-vat hypothesis, which imply the falsity of (nearly) all our external world beliefs, not merely one or two beliefs.²¹ So surely the logical possibility mentioned in the previous paragraph, that $D(p \wedge bb)$ and p are both true, will not apply to most beliefs p so long as the sceptical hypothesis is sufficiently radical?

This reply may seem tempting, and contains a germ of truth, but care is needed. Note firstly that even if we confine attention to radical sceptical hypotheses, it cannot *in general* be true that $D(p)$ and $p \Rightarrow q$ imply $\neg q$, where p and q are external world propositions; that is, the sceptical hypothesis $D(p)$ cannot imply the falsity of all the (worldly) consequences of p . For if this were true, we could immediately deduce that $D(p)$ is false for all p , i.e. we could know *a priori* that all sceptical hypotheses are false!²² Secondly, note that the compatibility of $D(p \wedge bb)$ and p is simply one way of illustrating the point that $D(p \wedge bb)$ does not distribute across conjunction (and so does not imply $D(p)$). The main argument for this point, given in section 2, is

that distributivity of D across conjunction, combined with D 's anti-factivity, implies that we can know *a priori* the falsity of certain sceptical hypotheses, which is implausible.

However – and this is the germ of truth – the sceptic is of course free to invent whatever sceptical hypothesis will best serve their purposes. So it is open to the sceptic to invoke the complex sceptical hypothesis $D(p \wedge bb) \wedge D(p)$ (“the demon is deceiving me into believing $p \wedge bb$ and the demon is deceiving me into believing p ”). This sceptical hypothesis obviously does imply $D(p)$, so this looks like a promising way of rendering the **CL** argument holism-proof while yielding the desired conclusion $\neg \text{JB}(p)$. But does it work?

5.3 Attempt 3: Conjoin-demonize-conjoin

Let the sceptical hypothesis be $D(p \wedge bb) \wedge D(p)$; again, this can be thought of as a rival to my conjunctive belief $p \wedge bb$. As before, the sceptic then tries to run the **UN** and **CL** arguments, with $p \wedge bb$ substituted for p and $D(p \wedge bb) \wedge D(p)$ for $D(p)$.

In the case of **UN**, this allows us to derive the sceptical conclusion $\neg \text{JB}(p \wedge bb)$ but no more, just as in Attempt 2 above. So we can confine attention to **CL**, which becomes:

CL1: if $\text{JB}(x)$, and it is known that $x \Rightarrow y$, then $\text{JB}(y)$

CL2***: $\neg \text{JB} \neg (D(p \wedge bb) \wedge D(p))$

therefore

C*: $\neg \text{JB}(p \wedge bb)$

As it stands, this is a valid argument for conclusion **C***. But the premises also entail the original sceptical conclusion **C**. For note that $p \Rightarrow \neg D(p) \Rightarrow \neg(D(p \wedge bb) \wedge D(p))$. Therefore, the sceptic can argue:

CL1: if $\text{JB}(x)$, and it is known that $x \Rightarrow y$, then $\text{JB}(y)$

CL2***: $\neg \text{JB} \neg(D(p \wedge bb) \wedge D(p))$

therefore

C: $\neg \text{JB}(p)$

This argument is valid, holism-proof, and yields the desired conclusion. So is the sceptic home and dry? No. For we need to ask whether premise **CL2*****, which says that I am not justified in believing the falsity of the complex sceptical hypothesis $D(p \wedge bb) \wedge D(p)$, is actually true.

Recall that the second premise of the **CL** argument is underpinned by the *Caro* principle. But it is far from clear that the *Caro* principle will give us **CL2*****. To see why, consider my belief p . Now the sceptical hypothesis $D(p \wedge bb) \wedge D(p)$ is certainly incompatible with p . But would this hypothesis, if true, explain my total evidence for p ? My total evidence for p is, by definition, my total evidence for $p \wedge bb$, i.e. the sensory experiences that form the evidential basis for $p \wedge bb$. But note that the first conjunct of this sceptical hypothesis, namely $D(p \wedge bb)$, *itself* explains those experiences. Thus the sceptical hypothesis $D(p \wedge bb) \wedge D(p)$ contains an irrelevant conjunct, so far as explaining my experiences concerned. And adding irrelevant conjuncts tends to destroy explanation.

Another way to see the point is this. In Attempt 2, when the sceptical hypothesis at issue was $D(p \wedge bb)$, the *Caro* principle could validly be invoked

to argue that I can't justifiably believe that the hypothesis is false; for the hypothesis, if true, would indeed explain the total evidence for p . But when the sceptical hypothesis is the more complex one $D(p \wedge bb) \wedge D(p)$, this reasoning breaks down, since it is not the hypothesis itself, but rather a proper part of it, that explains the evidence in question. In short, the *Caro* principle can establish **CL2**** but not **CL2*****.

Might it be replied that adding irrelevant conjuncts need not destroy explanation completely, but only weaken it? Perhaps. But in that case, the *Caro* principle would anyway need modifying. To see this, suppose we start with the sceptical hypothesis $D(p \wedge bb)$ and add the irrelevant conjunct $2 + 2 = 3$, to form the complex hypothesis $D(p \wedge bb) \wedge 2 + 2 = 3$. Now even if we grant that this complex hypothesis, and not merely its first conjunct, explains (to some extent) the evidence in question, it clearly cannot be right to conclude that I am unjustified in believing the negation of the hypothesis. For the negation of $D(p \wedge bb) \wedge 2 + 2 = 3$ is a necessary truth, which I surely am justified in believing! Thus if explanation is thought to be able to survive the addition of irrelevant conjuncts, the *Caro* principle would have to be restricted to cases where y is the *logically weakest* x -incompatible explanation of the total evidence for x .

In short, Attempt 3 looks tempting at the first sight, since by defining the sceptical hypothesis as $D(p \wedge bb) \wedge D(p)$, a valid closure-based argument for the original sceptical conclusion $\neg\text{JB}(p)$ can be given. But the soundness of the argument is open to question. There seems no reason why we should accept the second premise **CL2*****. In particular, the *Caro* principle does not establish **CL2*****.

5.4 Upshot

Let us summarize. To holism-proof the **UN** and **CL** arguments, the sceptic needs to consider background beliefs bb as well as the target belief p . This necessitates re-formulating the sceptical hypothesis, which can be done in three ways. The first way, “demonize-then-conjoin”, does not work, for the second premises of both sceptical arguments then become untenable. The second way, “conjoin-then-demonize” allows the sceptic to establish the second premises of both arguments, but the sceptical conclusion that results is the relatively unthreatening $\neg JB(p \wedge bb)$, rather than the original $\neg JB(p)$. There is no route to the latter conclusion; the route via **CL** fails since the D operator does not distribute across conjunction. The third way, “conjoin-demonize-conjoin”, at first sight seems to allow the original sceptical conclusion to be deduced via the **CL** argument, but closer investigation shows that the argument’s second premise becomes untenable, since the sceptical hypothesis contains an irrelevant conjunct.

6 Conclusion

Many epistemologists regard the **UN** and **CL** arguments as recipes for producing a sceptical conclusion about any of our external world beliefs. However as standardly formulated these arguments fall foul of evidential holism, for they presuppose a one-one evidential relation between episodes of sensory experience and beliefs about the world, which ignores the crucial role of background beliefs. The parallel with philosophy of science, where the problem that evidential holism causes for underdetermination-based sceptical arguments has long been noticed, makes this point explicit.

If holism is granted, this prompts the question of whether the **UN** and **CL** sceptical arguments can be rendered holism-proof, and if so how. In-

tuitively it seems that this should be possible, since the background beliefs are themselves about the external world so can be “demonized”, or brought within the scope of the sceptical hypothesis. By applying the sceptical arguments to beliefs of the form $p \wedge bb$ and taking the sceptical hypothesis to be $D(p \wedge bb)$ (as per Attempt 2), the sceptic can certainly deduce a conclusion of the form $\neg\text{JB}(p \wedge bb)$; however this sceptical conclusion is weaker than, and less threatening than, the original sceptical conclusion $\neg\text{JB}(p)$. There is no route to $\neg\text{JB}(p)$ on either argument.

It is perhaps somewhat surprising that holism serves to partially defuse the sceptical threat. One might have thought that holism would make epistemic justification harder, not easier, to achieve, since justifying a package deal seems harder than justifying a single belief.²³ However our aim here has not been to give a positive demonstration that our beliefs *are* justified, but rather to scrutinize the sceptic’s attempt to show that they are *not* justified. For this more modest aim, holism helps. For the sceptic’s reasoning involves constructing rival hypotheses that fit the evidence on which our everyday beliefs are based. So it is unsurprising that given evidential holism, the sceptic will at most be able to show that clusters of beliefs, not individual beliefs, are unjustified.

Our analysis has two main consequences. Firstly, it shows that the sceptical threat posed by the **UN** and **CL** arguments has been overestimated. Even granting the assumptions on which these arguments rest, discussed in section 2.3, the arguments fall short of their goal of showing that none of our worldly beliefs is justified, since they cannot be applied on a belief-by-belief basis.

Secondly and relatedly, it follows that the dialectical situation in regard to scepticism is not quite what has often been assumed. Many philosophers have thought that the only way to resist scepticism about the external world is to reject one or more of the assumptions on which **UN** and **CL** rest

(closure of justified belief under logical consequence, evidentialism about justification, failure of the “best explanation” response, sensory experience as the evidential basis for worldly beliefs). But this presumes that, granting these assumptions, the **UN** and **CL** arguments do succeed in establishing their intended conclusion, namely that none of our worldly beliefs is justified. Since holism-proof versions of these arguments only establish a weaker conclusion, it is not true that we must give up one of those assumptions on pain of conceding that none of our beliefs about the world is justified.

Notes

¹See in particular Yalçın (1992), Brueckner (1994), Cohen (1998), Vogel (2004), Pritchard (2005) and McCain (2013).

²Sceptical arguments that target justification are arguably more troubling. A sceptical argument to the effect that we don't *know* that p , even if successful, leaves open the possibility that we are perfectly justified in believing that p . So it is best to focus on justification-scepticism directly.

³Whether it is sensory experiences themselves, or beliefs about them, that stand in an evidential relationship to external world beliefs does not matter for present purposes. Also, the issue of whether sensory experiences do or do not have representational content may be set aside here.

⁴“Actual” serves to rigidify the definite description “the experiences on the basis of which I believe that p ”. Thus $D(p)$ is true at a world w iff the demon tricks me into believing p at w by causing sensory experiences indistinguishable from those on the basis of which I believe p in the actual world.

⁵Pritchard (2005), Brueckner (1994) and Cohen (1998) all employ formulations very similar to **UN1**.

⁶I have used x and y in place of McCain's p and q , in this quotation.

⁷See for example Sober (2008).

⁸It is sometimes thought that there is an alternative route to **CL2**, via the observation that if $D(p)$ were true, I would still believe it was false. However, this reasoning relies on a “sensitivity principle” for justification that is both open to counterexample and anyway conflicts with the closure principle that underpins premise **CL1**, as Comesaña and Klein (2019) observe.

⁹Thanks to an anonymous reviewer for this objection.

¹⁰See Pritchard (2005), Cohen (1998) and McCain (2013) for different views about the relation between **UN1** and **CL1**.

¹¹Pryor (2000) and Huemer (2001) argue that everyday perceptual beliefs are immediately justified by appearances.

¹²Vogel (1990) defends this line of argument. See Neta (2004) for a critical discussion

¹³Nozick (1981) and Dretske (1970) argue that knowledge is not closed under known logical implication; if correct, their arguments may apply to justification too.

¹⁴This move is made by direct realists and disjunctivists about perception (e.g. McDowell 2008); and in a different form by those who deny that sensory experience is “epis-

temically prior” to knowledge of the external world (e.g. Williams 1998), and those who reject the idea that our evidence is “the same” in two situations that are subjectively indistinguishable (e.g. Williamson 2000).

¹⁵See Laudan and Leplin (1991), Kukla (1996) and Stanford (2021) for discussion of this argument.

¹⁶This structural similarity is emphasized by Shogenji (2017), who explicitly describes Cartesian sceptical hypotheses as empirically equivalent alternatives to our everyday beliefs.

¹⁷See in particular Boyd (1973), Laudan and Leplin (1991) and Okasha (2002).

¹⁸The causal claim made by p^+ is arguably part of what I mean when I say that I can *see* that it’s sunny outside.

¹⁹For example, my background beliefs may include “there are no evil demons”.

²⁰Without this assumption, the sceptic’s attempt to holism-proof the **UN** and **CL** arguments would not get off the ground.

²¹Thorpe (2018) has recently emphasized that only radical sceptical scenarios are suitable for the sceptic’s purposes, as non-radical scenarios leave intact the justification of some of our empirical beliefs, which we can legitimately use to show that we are not in the non-radical scenario.

²²To see this, let x be any external world proposition that is logically independent of p . Note that $p \Rightarrow p \vee x$ and $p \Rightarrow p \vee \neg x$. So by the principle in question, $D(p) \Rightarrow \neg(p \vee x)$ and $D(p) \Rightarrow \neg(p \vee \neg x)$. But $\neg(p \vee x)$ and $\neg(p \vee \neg x)$ together imply both x and $\neg x$; therefore $D(p)$ is false.

²³Thanks to an anonymous reviewer for this point.

References

- Boyd, Richard (1973) “Realism, Underdetermination, and a Causal Theory of Evidence”, *Noûs* 7, 1, 1-12.
- Brueckner, Anthony (1994) “The Structure of the Skeptical Argument”, *Philosophy and Phenomenological Research* 54, 827–35.
- Cohen, Stewart (1998) “Two Kinds of Skeptical Argument”, *Philosophy and Phenomenological Research* 58, 1, 827–35.
- Comesaña, J. and Klein, P. (2019), “Skepticism”, *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), E. N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2019/entries/skepticism/>
- Dretske, Fred (1970) “Epistemic Operators”, *Journal of Philosophy* 67, 1007–1023.
- Huemer, Michael (2001) *Skepticism and the Veil of Perception*. Lanham, MD: Rowman and Littlefield.
- Kukla, André (1991) “Does Every Theory Have Empirically Equivalent Rivals?”, *Erkenntnis* 44, 137–166.
- Laudan, Larry and Leplin, Jarrett (1991) “Empirical Equivalence and Underdetermination”, *Journal of Philosophy* 88, 449–472.
- McCain, Kevin (2013) “Two Kinds of Skeptical Argument”, *Philosophy and Phenomenological Research* 164, 2, 289–300.
- McDowell, John (2008) “The Disjunctive Conception of Experience as Material for a Transcendental Argument”, in *Disjunctivism: Perception, Action, Knowledge*, F. Macpherson and A. Haddock (eds.), Oxford: Oxford University Press, 376–389.

- Neta, Ram (2004) “Skepticism, Abductivism, and the Explanatory Gap”, *Philosophical Perspectives* 14, 1, 296–325.
- Nozick, Robert (1981) *Philosophical Explanations*, Cambridge: Cambridge University Press.
- Okasha, Samir (2002) “Underdetermination, Holism and the Theory/Data Distinction”, *The Philosophical Quarterly* 52, 208, 303–319.
- Pritchard, Duncan (2005) “The Structure of Sceptical Arguments”, *The Philosophical Quarterly* 55, 218, 37–52.
- Pryor, James (2000) “The Skeptic and the Dogmatist”, *Nous* 34, 4, 517–549.
- Quine, Willard Van Orman (1951) “Two Dogmas of Empiricism”, *Philosophical Review* 60, 20–43.
- Quine, Willard Van Orman (1979) “Comments on Newton-Smith”, *Analysis* 39, 2, 66–67.
- Russell, Bertrand (1921) *The Analysis of Mind*, London: George Allen and Unwin.
- Salmon, Wesley (1984) *Scientific Explanation and the Causal Structure of the World*, Princeton: Princeton University Press.
- Shogenji, Tomoji, (2017) *Formal Epistemology and Cartesian Skepticism: In Defense of Belief in the Natural World*, London: Routledge.
- Sober, Elliott (2008) *Evidence and Evolution*, Cambridge: Cambridge University Press.
- Stanford, Kyle (2021) “Underdetermination of Scientific Theory”, *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.),

URL = <https://plato.stanford.edu/archives/win2021/entries/scientific-underdetermination/>.

Thorpe, Joshua (2018) “Closure Scepticism and the Vat Argument”, *Mind* 127, 507, 667–690.

Vogel, Jonathan (1990) “Cartesian Skepticism and Inference to the Best Explanation”, *Journal of Philosophy* 87, 11, 658–666.

Vogel, Jonathan (2004) “Varieties of Skepticism”, *Philosophy and Phenomenological Research* 68, 1, 1–37.

Williams, Michael (1998) “Epistemological Realism and the Basis of Scepticism”, *Mind* 97, 387, 415–439.

Williamson, Timothy (2000) *Knowledge and its Limits*, Oxford: Oxford University Press.

Yalçın, Ümit (1992) “Sceptical Arguments from Underdetermination”, *Philosophical Studies* 68, 1–34.