

The ‘Default View’ of Perceptual Reasons and ‘Closure-Based’ Sceptical Arguments

Genia Schönbaumsfeld

University of Southampton

Abstract

It is a commonly accepted assumption in contemporary epistemology that we need to find a solution to ‘closure-based’ sceptical arguments, and, hence, to the ‘scepticism or closure’ dilemma. In the present paper I argue that this is mistaken, since the closure principle does not, in fact, do real sceptical work. Rather, the decisive, scepticism-friendly moves are made before the closure principle is even brought into play. If we cannot avoid the sceptical conclusion, this is not due to closure’s holding it in place, but because we’ve already been persuaded to accept a certain conception of perceptual reasons, which both issues a standing invitation to radical scepticism and is endemic in the contemporary literature. Once the real villain of the piece is exposed, it will become clear that the closure principle has been cast in the role of scapegoat in this debate.

Keywords: closure principle, radical scepticism, Dretske, perceptual reasons, indistinguishability argument, reasons identity thesis.

I Introduction

The most popular way of motivating radical scepticism in the contemporary literature is by appeal to so-called ‘closure’-based sceptical arguments. The reason for this seems to be that such arguments generate the impression that the radical sceptical problem is natural and intuitive, since we need only to apply a logically compelling principle – such as the notion that knowledge is ‘closed’ under known entailment – to our ordinary perceptual beliefs in order to obtain the conclusion that unless we can rule out in advance that we are the victims of radical sceptical hypotheses, our ordinary beliefs are in jeopardy. ‘Closure’-based sceptical arguments tend to run as follows:

(BIV1) If I know I have two hands, then I know I’m not a brain-in-a-vat (BIV).

((BIV1) is motivated by the closure principle, ‘If S knows that p, and S competently deduces q from p thereby coming to believe that q on this basis while retaining her knowledge that p, then S knows that q.’^{1 2}). But, it seems,

(BIV2) I don’t know I’m not a BIV.

So,

(BIV3) I don’t know I have two hands.

In other words, if knowledge is ‘closed’ under known entailment, it appears that being able to rule out that one is a BIV has epistemic priority: the onus seems to be on the anti-sceptic to show that we do have perceptual knowledge of the world around us. If we cannot refute the idea that we might be the victims of radical sceptical scenarios, then, on pain of denying ‘closure’ – an intuitively plausible principle, rejection of which would seem to leave us at the mercy of deRose’s ‘abominable conjunctions’³ – our knowledge of the world is precarious.

So, unless a refutation of radical scepticism is available, closure-based sceptical arguments appear to confront us with two equally unpalatable options: either to embrace scepticism or to deny closure. What to choose? One famous response, offered by Fred

¹ I am here following Pritchard (2012). See also Hawthorne (2005).

² I am taking ‘being a BIV’ as shorthand for being the victim of a radical sceptical hypothesis.

³ Conjunctions of the sort ‘I know I have a hand, but I don’t know I’m not a handless BIV’ (see deRose (1995)).

Dretske (2005a)⁴, consists of biting the bullet and throwing out the closure principle, as this seems preferable to remaining stuck in the clutches of the radical sceptic: ‘The only way to preserve knowledge of homely truths, the truths everyone takes themselves to know, is...to abandon closure’ (Dretske 2005a: 18)⁵.

Although most contemporary epistemologists are not in favour of adopting such drastic measures, it is a commonly accepted assumption that we need to find a solution to closure-based sceptical arguments, and, hence, to the ‘scepticism or closure’ dilemma. In the present paper I argue that this is mistaken, since despite near-unanimity⁶ to the contrary in the current literature, the closure principle does not, in fact, do real sceptical work. Rather, the decisive, scepticism-friendly moves are made *before* the closure principle is even brought into play. If we cannot avoid the sceptical conclusion, this is not, I will argue, due to closure’s holding it in place, but because we’ve already been persuaded to accept a certain conception of perceptual reasons, which both issues a standing invitation to radical scepticism and is endemic in the contemporary literature. I am going to call this conception the ‘default view’.⁷ Once the real villain of the piece is exposed, it will become clear that the closure principle has been cast in the role of scapegoat in this debate⁸. So, even though my arguments are, in the first instance, aimed at Dretske, their moral has general application, and is relevant to anyone who accepts the ‘default view’ (that is to say, to the vast majority of contemporary epistemologists).

⁴ First formulated in his paper ‘Epistemic Operators’ (1970).

⁵ Compare Nozick (1981: 242): ‘If our notion of knowledge was as strong as we naturally tend to think (namely, closed under known entailment) then the sceptic would be right.’

⁶ Exceptions are Klein (2004), Floridi (2013), and, most recently Atkins and Nance (2014).

⁷ See footnote 21 for a (non-exhaustive) list of proponents of the view.

⁸ My argument also has implications for sceptical arguments employing the ‘underdetermination principle’. See footnote 33.

The strategy I adopt is as follows. First, I will briefly outline Dretske's reasons for rejecting closure.⁹ Then I will go on to show that Dretske's conception depends on some questionable – but widely shared – premises, and that it is those premises themselves, rather than an application of closure, that have the sceptical implications. Since the closure principle cannot provide us with any new empirical knowledge, it is the 'default view' of perceptual reasons that in fact causes the trouble. The point of this paper is not to refute radical scepticism, however, but rather to show that the problem does not lie where contemporary epistemologists commonly think it does.

II Dretske's Case against Closure

Dretske's first move consists of claiming that evidential transmission failure strongly suggests (but does not entail) the failure of closure (Dretske 2005a: 15). So if it can be shown that perception, our main route to knowledge of the world, does not transmit evidential warrant to all the known consequences of what is perceived, then this provides *prima facie* grounds to be sceptical of closure. Dretske argues that when we perceive that P, P has certain 'heavyweight' implications that cannot themselves be perceived to be so, and which we consequently have no reason to believe, even though these implications are among the known consequences of P (*ibid.*, 16). Dretske gives the following example, which is worth quoting in full:

⁹ I am using Dretske's position as a case study only, not because I think that his closure-denying 'anti-sceptical' strategy is the only game in town.

When Jimmy peeks into the cookie jar and, to his delight, sees that there are cookies there, his visual experience of the cookies, the evidential basis for his knowledge that there are cookies there, is not evidence, not a reason to believe, that there is a physical reality independent of Jimmy's mind. Jimmy's experience of the cookies may be good reason to believe there are cookies in the jar, but it is not a good reason to believe that idealism is false. And it is not a good reason to believe that idealism is false even if Jimmy understands that cookies are mind-independent objects and that, therefore, what he sees to be the case (that there are cookies in the jar) implies that idealism is false. Looking in the cookie jar may be a way of finding out whether there are any cookies there, but it isn't – no more than kicking rocks – a way of refuting Bishop Berkeley (*ibid.*, 15).

In other words, the 'heavyweight' implication of Jimmy's seeing that there are cookies in the jar is that idealism is false (or that Jimmy is not a BIV), but this implication is not something that Jimmy can 'perceive to be so', and, hence, have any evidential warrant for. So, perception, according to Dretske, never gives one reason to believe that the 'heavyweight' implications of ordinary perceptual claims are true, even though one knows that they figure among the known consequences of these claims.

The upshot of this is that if Dretske is right about perceptual evidential transmission, and we accept the closure principle, then a 'closure-based' sceptical argument can run¹⁰. For if closure is endorsed, then, if Jimmy knows that there are cookies in the jar, he can deduce, and thereby come to know, the 'heavyweight' implication that he is not a BIV (or that idealism is false), since being a BIV is incompatible with knowing that there are cookies in

¹⁰ The following is my reconstruction of Dretske's argument, not a verbatim account.

the jar (given that if Jimmy were a BIV, he would not be perceiving actual cookies, but merely 'vat' cookies). So:

(BIV 1) If Jimmy knows that there are cookies in the jar, then Jimmy knows that he is not a BIV (or that idealism is false).

But since Dretske has argued that seeing that there are cookies in the jar does *not* give Jimmy good reason to believe the 'heavyweight' implication that he is not a BIV:

(BIV 2) Jimmy does not know that he is not a BIV (or that idealism is false).

So, by *modus tollens*:

(BIV 3) Jimmy does not know that there are cookies in the jar.

Hence, it now seems that unless we deny closure – i.e. unless we deny that we need to know their 'heavyweight' implications in order to know the ordinary things we think we know (such as that there are cookies in the jar) – we have to accept the sceptical conclusion. In the face of this, Dretske contends, it is surely better to maintain that there are some things, namely, the 'heavyweight' implications, that we needn't know, despite the fact that we know that our ordinary knowledge depends on their truth (*ibid.*, 17). In this way the 'appearances'

can be saved without succumbing to scepticism. In the remainder of his paper, Dretske attempts to discredit possible alternatives to his view – attributor contextualism¹¹, and what he calls ‘verbal hocus pocus’ (*ibid.*, 23) – in order to show that denying closure is really the only game in town.

Contextualism is the view that knowledge attributions are sensitive to justificational context, so it is possible to know something in, say, an ‘ordinary’ context, but not to know the same thing in a ‘sceptical’ context, for example, where the stakes are higher. In other words, advocates of contextualism can claim that it is possible to know that there are cookies in the jar in ordinary justificational contexts where no sceptical error-possibilities have been raised, while, at the same time, being able to maintain that these things can no longer be known in contexts where sceptical error-possibilities are under consideration. In this way, contextualism can retain closure without selling out to scepticism (since it is still possible to know things in ordinary contexts).

Dretske is not impressed by this manoeuvre, however, as he believes that it is a consequence of this view that philosophers who spend time worrying about heavyweight implications ‘are the most ignorant people in the world. Not only don’t they know these heavyweight implications (maybe no one does), they don’t (like everyone else) know the things (that there are cookies in the jar, zebras in the pen) that imply them’ (*ibid.*, 19)¹². This, Dretske claims, ‘is pretty bizarre’ – much ‘*more* bizarre’ than the attempt to deny closure (*ibid.*).

The view that Dretske labels ‘verbal hocus pocus’, on the other hand, is the contention that we can come to know ‘heavyweight’ implications without having reasons to believe them – without, for example, being able to eliminate the error-possibility that one is a BIV

¹¹ See, for example, Cohen (1999, 2005); deRose (1995).

¹² This may well be unfair to contextualists, but I am not going to engage in this debate here.

who is being ‘fed’ cookie-experiences (*ibid.*, 23). Dretske admits that dismissing this notion as hokey requires the assumption that the ‘appearance of cookies in a jar, the sort of condition that prompts one to say one can see there are cookies in the jar, is not itself a reason to believe that the experience of the cookies is not misleading or delusory in some way’ (*ibid.*, 24). Granted this assumption, Dretske claims it is ‘abominable’ nevertheless to insist that we know we are not being deceived because we know that not being deceived is implied by what our experience normally leads us to believe (*ibid.*)¹³.

III The Case Against Dretske

It is clear that the main thrust of Dretske’s case against closure is centred around the thought that seeing that there are cookies in the jar is not a reason to believe that there is a physical reality outside of one’s mind (that there really *are* cookies in the jar¹⁴), since, if one were a BIV, things would look exactly the same as if one weren’t: ‘If your reasons for believing P are such that you *might* have them when P is false, then they aren’t good enough to *know* that P is true’ (Dretske (2005b: 44))¹⁵. The question is, what reasons do we have for accepting this claim and are they strong enough to motivate a rejection of closure?

So, let’s dig a little deeper. What considerations speak in favour of endorsing the view that your reasons for believing P – say, seeing that there are cookies in the jar – are reasons

¹³ I have not, in this section, endorsed any of Dretske’s claims or assessed their accuracy – I am merely reporting his views.

¹⁴ I take it that if cookie-experiences do not give one reason to believe that there is a mind-independent reality, then cookie-experiences do not give one reason to believe that there is a *mind-independent, physical* (i.e. real) cookie in front of one either (otherwise (‘good case’) cookie-experiences would give one reason to believe that there is a mind-independent reality).

¹⁵ In his (1971) Dretske puts this point in terms of ‘conclusive reasons’: ‘If S knows that P on the basis of evidence E, S would not have E unless P were true’. Since Dretske believes that such conclusive reasons have to rule out sceptical alternatives, perceptual reasons, on Dretske’s account, are never conclusive. Hence, one can have conclusive reasons for *believing* P, but not for *knowing* P (compare Dretske (2010)).

that you would have even when P is false? For, *prima facie*, this is not as obvious as Dretske seems to think, since one might believe instead that in cases where P is false, it only *appears as if* one had reasons to believe that P, not that, *in actual fact*, one has reasons to believe that P. That is to say, one might think instead that in cases where P is true, we have reason to believe that P, while in cases where P is false, we do not have reason to believe that P – it only *seems* that we do¹⁶. Consequently, *pace* Dretske, there is no *prima facie* reason to concede that the ‘good case’ and the ‘bad case’ are ‘symmetrical’ in the sense that they provide us with exactly the same perceptual reasons.

Why does Dretske not consider this possibility? Perhaps because he is implicitly relying on an argument from subjective indistinguishability: unless I can ‘tell’ merely by ‘inspecting’ or ‘consulting’ my immediate cookie-experience that it is in fact an experience of a real cookie and not, say, of a ‘vat-induced’ one, I cannot know that I am experiencing a *real* cookie. Since it seems impossible so to distinguish, cookie-experiences do not give one an evidential basis for knowing that there are (real) cookies there.

Despite appearances to the contrary, however, arguments from subjective indistinguishability are fallacious. A reconstruction of the argument can help us to see this:

The Indistinguishability Argument (IA)

(P1) In the bad case (i.e. where it is false that there are cookies in the jar), cookie-experiences do not give one reason to believe that there are cookies in the jar (since there are no cookies in the jar).

¹⁶ So, in such cases one might be blamelessly taking oneself to be in possession of a reason, even though one is not. But blamelessness alone does not suffice to turn an appearance (of a reason) into reality.

(P2) The good and the bad cases are subjectively indistinguishable.

(C1) So, in the good case, cookie-experiences do not give one reason to believe that there are cookies in the jar.¹⁷

Not only is this argument not valid, it also wreaks havoc on Dretske's case against closure (as we shall see below).

To see why the argument is invalid, consider the following three consistent claims:

(Q1) In the bad case, my cookie-experiences are not of cookies.

(Q2) The good and the bad cases are subjectively indistinguishable.

(Q3) In the good case, my cookie-experiences are of cookies.

If (Q1), (Q2) and (Q3) are consistent, however, then (Q2) does not entail (C1):

(Q1) In the bad case, my cookie-experiences are not of cookies.

¹⁷ Compare Pritchard ((2008), (2009), (2012)).

(Q2) The good and the bad cases are subjectively indistinguishable.

(C1) So, in the good case, my cookie-experiences are not of cookies.

In this argument, it is quite easy to see that phenomenological (subjective) indistinguishability does not imply non-distinctness. For, surely, nobody would want to claim that, in the good case, my cookie-experiences are not of cookies, even though it is not possible subjectively to distinguish between the good and the bad case. So, the fact that I might not, at any one time, be able subjectively to distinguish between a real cookie and a 'vat-induced' one, does not imply that my experiences of them are not distinct. But, if so, then why should we accept the thought that the reasons the two different scenarios supply must nevertheless be the same (call this the 'Reasons Identity Thesis' – RIT). That is to say, why should we accept that *even* the good case cannot provide us with better reasons than the bad case? For, one might think that in the good case, my cookie-experiences can be factive¹⁸, while, in the bad case, they cannot. So, if the good and the bad cases themselves are asymmetrical, why should this not also apply to the reasons that these different scenarios can provide us with? In other words, one might maintain that if one is in the good case, where one's cookie-experiences can be factive, one can come to know on this basis that there are cookies in the jar, although one cannot come to know this in the corresponding bad case, where it would only seem to one as if one had good reason to believe that there are cookies in

¹⁸ By a 'factive' cookie-experience I mean an experience whose presence entails that there is a cookie there.

the jar. To rule this option out *ab initio*, or by continuing to insist on subjective indistinguishability, would be to beg the question against one's opponent¹⁹.

If this is right, then we have reason to be sceptical of Dretske's claims about evidential transmission failure, and, hence, of his grounds for rejecting closure. For if, in the good case, my perceptual experience of cookies can be factive, then we need not accept Dretske's argument that we can *never* 'perceive heavyweight implications to be so'. One might, instead, reason as follows: If I am in the good case, and I can come to know, in virtue of my *factive* cookie-experience, that I am seeing a cookie, then I can also come to know – in virtue of my *factive* cookie-experience – that I am not seeing a 'vat-cookie', as seeing a 'vat-cookie' is incompatible with seeing a *cookie*²⁰. But coming to know that I am not seeing a 'vat-cookie' is not, *pace* Dretske, an 'extra', as it were 'heavyweight' experience that I might have 'on top of' seeing that there is a cookie there. Rather, if I am in the good case, and RIT is rejected, then seeing that there is a cookie there is *already*, all by itself, a suitably 'heavyweight' proposition. That is to say, if my cookie-experience is factive, then it is already an experience, as it were, of a 'heavyweight' cookie, not, *pace* Dretske, an experience of something that could *either* be the experience of a 'heavyweight' cookie *or* the experience of a 'vat-cookie' (since, on such a conception, no cookie-experience could ever be factive – that is to say, entail the presence of an actual cookie, not the mere appearance of one)²¹.

That Dretske thinks otherwise – i.e. that Dretske believes that seeing that there is a cookie in the jar cannot 'inform' one about the 'heavyweight implication' that there is not a

¹⁹ Compare Williams (1996: 79-88) and Millar (2011).

²⁰ I am here leaving it open whether we can know that we are in this privileged situation, I am merely arguing that this is a possible alternative to Dretske's conception and that he has not done enough to rule it out.

²¹ This conception of (non-conclusive) perceptual reasons seems to be the default position in contemporary epistemology. According to this 'default view', 'cookie-experiences' are always experiences *as of there being cookies* (and hence non-factive), rather than experiences of 'real' or 'heavyweight' cookies (what I am calling factive cookie-experiences). This view is endorsed, *inter alia*, by Bennett (1971), Burge (2003), Conee (2007), Coliva (2012), McGinn (1984), Millar (1991), Nagel (1986), Pollock (1974), Pryor (2000), Stroud (1984), White (2014), Wright (2002, 2014).

‘vat-cookie’ in the jar and, consequently, that we need to deny closure in order to save ourselves from scepticism – is due to his prior commitment to the truth of the indistinguishability argument (IA) and his concomitant endorsement of the Reasons Identity Thesis (RIT), something that comes out clearly in the following remark. Dretske says:

Despite knowing that cookies are objective (mind-independent) objects, I can see (roughly: tell by looking) that there are cookies in the jar without being able to see, without being able to tell by looking, that there are mind-independent objects...Maybe one has to know there are physical objects in order to see that there are cookies in the jar..., but one surely isn’t claiming to see that there are physical objects in claiming to see there are cookies in the jar. After all, hallucinatory cookies ‘in’ hallucinatory jars can look exactly like real cookies in real jars. So one cannot, not by vision alone, distinguish real cookies from mental figments. One cannot *see* that the world really is the way it visually appears to be (Dretske (2005a: 14)).

In other words, Dretske claims that unless one can perceptually (visually) distinguish real cookies from mental figments (or ‘vat-cookies’), seeing that there are cookies in the jar is *never* a reason to believe that there are real (‘heavyweight’) cookies in the jar: ‘Perceptual reasons – the sense-experiences on which we base everyday perceptual judgements – do not transmit their evidential force to all the known consequences of the judgments they warrant’ (Dretske (2005a: 15)). That is to say, on the ‘default view’ that Dretske is espousing, even in the good case our perceptual reasons can never be better than in the bad case, since how things appear to us (the information we receive from our sense-experiences) is the same, whether we are deluded, hallucinating or envatted, or not. Dretske couldn’t make it clearer

that he is endorsing the Reasons Identity Thesis and that he is basing it on the conclusion of IA.

Since we have already seen, however, that IA is fallacious, Dretske is mistaken to think that phenomenological indistinguishability alone is enough to motivate RIT. For even if one cannot subjectively distinguish a 'vat'-cookie from a real cookie, it does not follow that therefore cookie-experiences (in general) can never be factive – i.e. that seeing that there is a cookie there can *never* (even in the good case) inform one of the presence of *real* (physical) cookies. Given that Dretske's case against closure assumes the truth of this conclusion, however, it seems that Dretske's case has just collapsed.

Let me reiterate exactly what the problem is here. The bone of contention is not whether, in any given scenario, one is able to know that one is in the good case, as for the purposes of this paper I am going to remain agnostic on this question. Rather, the issue is whether, *were one to be in the good case*, one's perceptual evidence would transmit to what Dretske calls 'heavyweight' implications. Anyone who accepts RIT contends that even if I were to find myself in the good case, my perceptual evidential grounds would not be factive, and, hence would not entail the presence of an actual physical object (a real cookie, hand, or what have you). So, on this view, one's perceptual experiences *as a class*, even in the good case, fall short of providing knowledge of an independent physical world. It is this view, I will argue, not an endorsement of closure, that has the sceptical implications, and this view that we consequently need to reject if any progress is to be made against the radical sceptic.

IV Scepticism and the 'Default View'

Since we have seen that, according to the ‘default view’, even the presence of a ‘good-case’ cookie-experience never provides a conclusive reason to believe, and hence to come to know, that there is a real physical object in front of one (a cookie, say), one wonders what it would take for experiences of *physical objects* to be possible. For if to see a cookie is never to see a mind-independent, physical object, as Dretske claims (in the quote cited above), then what, on Dretske’s conception, would have to be the case in order for physical object perception to be possible? That perceptual experience come along with ‘experiential markers’ that would allow one, as it were ‘from the inside’, to distinguish between veridical perception and illusion²², or that perceptual experiences be ‘self-identifying’ – capable, somehow, of ‘telling me’ that they are not ‘vat-induced’ experiences? Not only is such a thing in principle impossible, it is a barely coherent notion. So one shouldn’t be surprised that if *this* (impossible) criterion is a necessary condition for the possibility of physical object perception, then physical object perception is ruled out from the start. But now it is beginning to look as if Dretske is just a sceptic in disguise, for he seems to be denying that physical object perception is ever possible, and that, surely, is a sceptical conclusion if ever there was one.

Even if one were to grant Dretske this contentious, scepticism-inducing point for the sake of argument, however, the fact that perceptual experience does not come with such ‘markers’ is not, as we have seen, sufficient to rule out the thought that *veridical* (good case) perception might nevertheless be factive. To claim otherwise is just to reiterate the fallacious argument that phenomenological indistinguishability implies RIT. In other words, we cannot argue that since, in the bad case, we are not in touch with the facts, and we cannot subjectively distinguish between the good and the bad case, therefore, in the good case, we are not in touch with the facts, as this would be a version of the invalid argument discussed

²² Dretske, in fact, seems to have something like this in mind; see Dretske (2010).

above. Consequently, we have, so far, been given no good reason for Dretske's claim that seeing a cookie in front of one is *never* a reason to believe that there is a *real* (physical) cookie in front of one.

The notion, moreover, that one doesn't (ever) have to see real cookies in order to know that one is having *cookie*-experiences – as opposed to experiences of anything you please – implies that cookie-experiences are completely detachable from cookies: one can, on this view, know that one is having a cookie-experience without thereby (ever) coming to know that one is perceiving a real (mind-independent) cookie. If we are willing to entertain the idea, as Dretske clearly is, that this may *always* be the case (i.e. that cookie-experiences *never* give one knowledge of the presence of sweet physical objects), then this implies that the only experiences available to one are experiences *as of there being cookies in front of one* – that is to say, experiences that could always *either* be the experience of a real cookie (or real hand etc.) *or* the experience of a 'vat'-cookie.²³

If this represents Dretske's conception, however, then it spoils his point about denying closure. For it now seems that, strictly speaking, we can never know that there are cookies in the jar, if by 'knowing that there are cookies in the jar' we mean what we normally mean, namely, 'seeing that there are real (physical) cookies in the jar' (what one might call, using Dretske's terminology, seeing 'heavyweight' cookies). Rather, all we can know, on his conception, is that we are having experiences as of there being cookies there – which of course leaves it open whether these experiences are (ever) of actual cookies or not. This is a result that the sceptic might happily accept: as long as we claim to know only about how we experience the world to be, rather than what the world is actually like, the sceptic has no quarrel with us. But, if this is so, then why does Dretske think that he has to deny 'closure' in

²³ This epistemic view also seems to have a rather unwelcome metaphysical implication, namely the inaccessibility to direct perception of mind-independent physical objects.

order to get himself this result? It seems that he is asking us to pay an extremely high price – to discard the plausible notion that we can extend our knowledge through deduction – for something we already have anyway (namely, knowledge of how things appear)²⁴.

What has gone wrong here? The problem lies with Dretske's equivocating on the meaning of the antecedent in (BIV1) and his consequent confusion about the role the closure principle plays in debates about radical scepticism. The reason why I cannot, for example, reason from 'I see that there is a cookie in front of me' to the presence of a 'heavyweight' cookie is *not* because closure fails (as Dretske claims), but rather because Dretske has already accepted the view that, even in the good case, cookie-experiences always fall short of providing knowledge of real (physical) cookies (i.e. that he has already accepted RIT – the 'default view'). Hence, all I can know, on this conception, is that I am having experiences as of there being cookies in front of me, but such non-factive experiences are of course never sufficient to ground knowledge-claims about 'heavyweight' cookies (about the fact that there are real cookies there). But, if this is so, then (BIV1), on Dretske's conception, should really read like this:

(BIV1)* If I know that I am having an experience as of there being a cookie in front of me, then I know I'm not a BIV.

Since one cannot deduce that one is not a BIV from the presence of a non-factive cookie-experience (an experience as of there being a cookie there), however, (BIV1)* is false and cannot serve as a premise in a closure-based sceptical argument. Consequently, once Dretske's equivocation is exposed, it becomes apparent that no closure-based sceptical

²⁴ For a similar point see Pritchard (2007: 41).

argument can get off the ground and hence that a denial of closure is entirely redundant. Given that Dretske is already *starting* with the claim that one can only have knowledge of cookie-appearances (non-factive cookie-experiences; experiences as of there being cookies there), never of the (physical) cookies themselves, and one cannot deduce something from a proposition that isn't entailed by it, the denial of closure is a red herring. For to deny closure requires denying that one can 'deduce' (and thereby come to know) that one is not a BIV from the presence of a 'heavyweight' cookie (from a *factive* cookie-experience), as only the presence of a factive cookie-experience entails the non-presence of a 'vat-cookie'. But that is not what Dretske is denying²⁵. Dretske is denying that one can deduce that one is not a BIV from a *non-factive* cookie-experience (a cookie-appearance), and this is both true and entirely consistent with the closure principle²⁶.

Conversely, if we reject the 'default view' and with it the claim that cookie-experiences can *never* – even in the best possible case – give one knowledge of 'heavyweight' cookies, then, at least in principle, physical cookie perception is possible. But, if so, then were one to find oneself in the good case where one's cookie-experience would be factive, and one could come to know this, then one could also come to know on this perceptual basis that one is not seeing a 'vat-cookie', as seeing a 'vat-cookie' is incompatible with seeing a (real) cookie. Consequently, regardless of whether or not one actually makes this 'deduction', the fact that one sees a real cookie can, in the good case, put one in a

²⁵ The only reason why this is not obvious is that talk of 'cookie-experiences' is neutral between 'factive' and 'non-factive' cookie-experiences, and so Dretske can equivocate between them.

²⁶ Compare Williams (1996: 190-1): 'In fact, if the sceptic is allowed this account [of non-conclusive reasons], the apparent failure of closure is an illusion. Those who deny closure claim that, in an ordinary context, I can know that there is a *hand* in front of me even if I do not know that I am not merely dreaming that there is. But the sceptic will reply that this is not really a case in which I know that P, know that P entails not-Q and yet fail to know that not-Q. Rather, it is a case in which I know-for-all-practical purposes that P, and so would ordinarily quite properly be *said* to know that P, even though, strictly speaking, I fail to know that P precisely because I fail to know that not-Q.'

position to come to know that one is, and in this much, that one is not perceiving a ‘vat-cookie’ also²⁷.

It is important to note here, however, that, if one were in such a good case scenario, the fact that one would not be seeing a ‘vat-cookie’ would follow from the fact that one sees a (real) cookie, *not* – as one might mistakenly suppose – from the fact that one has already accepted the closure principle.²⁸ Hence, what would be doing the anti-sceptical work here would be one’s factive cookie-experience, *not* the fact that knowing that there is a real cookie there implies that there is not a vat-cookie in front of one. Conversely, if one can never know that one is perceiving a real cookie, because all one can ever know is that one is having an experience as of there being a cookie in front of one, then what is doing the sceptical work is the general non-factivity of cookie-experiences, not the closure principle.

An example from a different context may help to illustrate this point. Consider, for instance, the following piece of reasoning that a detective might engage in:

(P1) If I know that Pierre is the murderer, then I know that Pierre was at the scene of the crime.

(P2) I know that Pierre was not at the scene of the crime.

(C1) I know that Pierre is not the murderer.

²⁷ I am merely making a logical point here, not committing myself on the question whether we can come to know that we are in the good case.

²⁸ So, I can only refute idealism by kicking a stone, say, if it is already granted that ‘stone-experiences’ are factive. Since the idealist will not grant this without further argument, however, such attempts remain question-begging (compare Atkins and Nance (2014) who argue that closure-based sceptical arguments are question begging in the same way that Moore’s ‘proof’ is).

It is an analytic truth that if someone is a murderer and has killed someone, then this person was at the crime scene²⁹. Hence (P1) is true on purely *a priori* logical grounds. I don't need to know the actual facts of a case in order to know that if someone is a murderer, then s(he) was at the scene of the crime. In this respect, (P1) is like (BIV1) in the 'closure-based' sceptical argument – it is a conditional that informs one about a logical relationship that holds between the antecedent and the consequent. Hence, (P1) and (BIV1) are both true regardless of what the (empirical) facts are. Just as I cannot know that I am perceiving a real object (a cookie, a hand or whatever), and not know that I am not systematically deluded, for example, so I cannot know that someone is the murderer and not know that this person was at the scene of the crime.

Now suppose that the detective says the following to his commanding officer: 'I know that Pierre is not the murderer.' His superior understandably goes on to ask the detective what his evidence for this claim is. Imagine the detective responds, 'If I know that Pierre is the murderer, then I know that Pierre was at the scene of the crime.' At this point, the commanding officer becomes impatient and thunders, 'Yes, and what of it?'

Why is this relevant? It is relevant because it shows that being in possession of 'deductive knowledge' that 'if I know that P, and I know that P entails Q, then I know that Q' tells us nothing whatsoever about *whether* I know that P. Since, however, I only come into possession of actual knowledge that Q (rather than knowledge merely of the *conditional relation* that *if* I know P, and I know that P implies Q, then I can deduce Q from P and thereby come to know Q), if I already am in possession of knowledge that P, what is, as it were, doing all the 'knowledge-bestowing' work is the fact that I know that P, *not the fact that my knowing that P enables me to deduce that Q*.

²⁹ I am here excluding the more 'exotic' forms of murder 'at a distance', or by proxy, where there are no 'crime scenes' in the relevant sense.

Analogously, the detective can only know that Pierre is not the murderer if he knows that Pierre was not at the scene of the crime (P2). But he can only know *that* if he has independent, empirical evidence for Pierre's not having been at the scene of the crime. He cannot 'deduce' that Pierre is not the murderer from the fact that Pierre was not at the scene of the crime unless he *already* knows that Pierre was not at the scene of the crime (and has evidence for it). Similarly, I could only come to know that I am not a BIV by deducing this from an ordinary perceptual claim, say, if I *already* know what I am perceiving and have factive epistemic support for it. But if I do have such support, then I can also deduce, and thereby come to believe, that I am not a BIV, since coming to know, on the basis of making such a deduction, that one is not a BIV is not a matter of coming to know a new, empirical fact by *a priori* means, but rather becoming aware of a logical implication of what one (already) knows³⁰. In this respect, the closure principle cannot furnish one with new empirical knowledge, but only informs one of what follows logically from what one already knows. Consequently, competent deduction cannot give one any anti-sceptical knowledge that one did not already possess before (one made the deduction) – it can only make one aware of the fact that the actual knowledge one possesses has anti-sceptical implications.

If this is right, then it turns out that *no* closure-based sceptical arguments are in fact possible³¹. For they either, as in Dretske's case, already fall at the first hurdle (i.e. once we become aware of Dretske's equivocation, we see that (BIV1)* cannot serve as a premise in a closure-based sceptical argument), or it is possible to deny (BIV2). In other words, we can either never know more than that we are having (non-factive) cookie-experiences, in which case we cannot deduce (and thereby come to know) that we are not perceiving 'vat-cookies' (or that we are not BIVs), since the presence of (non-factive) cookie-experiences does not

³⁰ Namely, that one can't be deluded, hallucinating or a BIV and know what one believes. (The 'BIV' scenario is merely shorthand for being in any kind of bad case.) Compare Millar (2014: 144).

³¹ This is not to say, of course, that *other* sceptical arguments aren't possible.

entail the non-presence of ‘vat-cookies’; or, if we reject RIT and accept that cookie-experiences can, in the good case at least, be factive, then, we can also, in principle, come to know on this basis that we are not perceiving a vat-cookie (or that we are not BIVs), so (BIV2) will come out as false. On neither horn of this dilemma is a closure-based sceptical argument possible: We either start with the scepticism-friendly assumption that all cookie-experiences as a class are non-factive, in which case we are already well on the way towards motivating global scepticism³², or we grant that cookie-experiences can, in the good case at least, potentially give us knowledge of real cookies, and, so, if we could come to know that we are in this privileged situation – a question that I have not pursued in this paper – then this would enable us to deny (BIV2). But whichever horn of this dilemma we grasp, the important point to appreciate is that neither of them is motivated by an acceptance or a denial of closure. It is rather completely independent considerations that force the choice – namely, whether or not good-case cookie-experiences can be factive, something that any proponent of RIT and the ‘default view’³³ would deny, and whether, if we reject this view, we can come to know that we are in the good case when we are.

Consequently, the closure principle itself is entirely blameless in the whole affair, since it operates in a purely formal, non-material manner. That is to say, it is not closure that can tell one anything ‘substantial’, such as, for example, whether one can know that one is in the good case or not. This question has to be settled in an entirely different way by, for instance, determining how good one’s grounds are for making the claim to know; whether

³² For if, say, one cannot know that one is perceiving a real cookie, a real hand, or is writing on a real computer, how can one really know anything?

³³ The ‘default view’ is also responsible for motivating radical scepticism from the underdetermination principle (UP): ‘If S ’s evidence for ϕ does not favour ϕ over an incompatible hypothesis ψ , then S is not justified in believing ϕ and rejecting ψ ’ (see Brueckner (1994) and (2010) for further discussion). An endorsement of this principle, however, only leads to scepticism on the assumption that my perceptual evidence is the same whether I am a BIV or not, an assumption that very clearly relies on IA. Hence, in this case too, IA is the culprit, not UP (for a good exposition of how UP (on the assumption that IA is true) leads to radical scepticism, see Yalcin (1992); for more discussion on the relation between UP and closure, see Pritchard (2005).

there is, say, any reason to suppose that there might be ‘cookie-façades’ in the vicinity, that someone might have slipped a hallucinogenic drug in one’s drink, that the lighting conditions are bad etc. In the absence, however, of the acceptability of the RIT-driven thought that perceptual evidence can *never* – even in the good case – be factive, there is no reason to grant Dretske the move that the (factive) good case is impossible and that perceptual evidential warrant therefore never transmits to ‘heavyweight’ implications.

In other words, no real (material) work is being done by (BIV1) at all, as it merely informs one of a logical relationship between two propositions which holds whether or not it is actually the case that one knows that P (the antecedent in (BIV1)). But if one *does* know that P, and one knows that P entails Q, then one can also come to know that Q. In other words, if, *pace* the ‘default view’, it is granted that in the good case one can, in principle, come to know that there is a real cookie there, then one can also come to know on this basis that there is not a ‘vat-cookie’ there. Conversely, if one can never know that there are ‘real’ cookies there, since cookie-experiences *as a class* always fall short of providing knowledge of the facts, then one can never deduce from their presence that one is not perceiving ‘vat-cookies’ either.

Dretske’s claim, therefore, that closure is not an acceptable epistemological principle since ‘you can know P is true and know $P \rightarrow Q$ and still not be in a position to know, even with the help of deductive reasoning, that Q is true’ (Dretske 2010: 134) is false. Closure is a perfectly acceptable epistemological principle. If you don’t know that Q is true, this is because you never really knew P in the first place.

V Some Objections

Does anything I said in the previous section commit me to what Dretske calls ‘verbal hocus-pocus’? That is to say, do my arguments imply that we can have knowledge of ‘heavyweight’ implications without having any reasons for them? No. Recall that Dretske believes that seeing that there is a cookie in front of one does not give one reason to believe the ‘heavyweight’ implication that one is not perceiving a ‘vat-cookie’, as it is impossible to ‘perceive’ that one is not perceiving a ‘vat-cookie’ (or that one is not a BIV). If one cannot ‘perceive’ this, however, then how can one have any reason for ruling out this error-possibility, and, hence, for coming to know that one is perceiving a real cookie and not a ‘vat-cookie’?

Dretske is right, of course, that one cannot ‘perceive’ that one is not perceiving a ‘vat-cookie’, since all one can perceive is a *cookie*. But unless one has already ruled in advance that perceiving a cookie can *never* give one knowledge of ‘heavyweight’ cookies, perceiving a cookie can, in the good case, be sufficient to give one knowledge of cookies even if one can’t subjectively distinguish between cookies and ‘vat-cookies’³⁴. Consequently, if RIT is rejected because its acceptance depends on the conclusion of the fallacious IA, then the fact that there is a cookie there is, in the good case, one’s reason for believing that one is not perceiving a ‘vat-cookie’, and, hence, *pace* Dretske, one doesn’t need an *additional* reason to believe the ‘heavyweight’ implication that one isn’t perceiving a ‘vat-cookie’.

In other words, if I am in the good case, no further evidence is required to license the move from ‘I am having a (factive) cookie-experience’ to ‘I see that there is a (“heavyweight”) cookie in front of me’, although in the corresponding bad case, additional evidence may be necessary (e.g. if I find myself in ‘cookie-façade’ land). Consequently,

³⁴ Compare Klein (2004), who makes a similar point about zebras and cleverly disguised mules; see also Pritchard (2012), parts II and III.

contra Dretske, we have no reason to believe that there is a *general* failure of evidential transmission to what he calls ‘heavyweight’ implications (only that this may sometimes be the case). Things will only seem otherwise if we grant Dretske his assumption that cookie-experiences can never – even in the good case – be factive (i.e. if we grant him RIT). Since this assumption presupposes the validity of IA, there is no obligation to make this concession.

Similarly, nothing much follows from the fact that one cannot ‘perceive’ that one is not a BIV either. Given that the BIV hypothesis is constructed in such a way that it is *logically* – not empirically – impossible for one to have any empirical evidence for it, it is impossible either to ‘perceive’ or not to ‘perceive’ that one is a BIV. But, if so, then appeals to perception are just irrelevant in this context. That is to say, no one entertains the BIV hypothesis because there is a presumption in its favour. In this much it is nothing more than a logical possibility, but logical possibilities alone are not reasons to believe something (or, indeed, to take them seriously). Rather, if something is not logically possible, then it need not even be investigated, but this does not imply that logical conceivability counts, by itself, as a reason that speaks in favour of a given scenario. Klein provides a good example to illustrate this point. In order to have sufficient evidential support for the claim that a particular person that we see is Publius, for instance, we don’t first need to be able to rule out an incredulous bystander’s idea that it might not be Publius, but rather Magicus’s dog, whom Magicus has transformed into an exact duplicate of Publius, unless we already have independent grounds for thinking that Magicus has that ‘ability, an opportunity and a motive’ (Klein (2004: 174)). If this is right, then we analogously don’t need to accept Dretske’s claim that we need to be able to rule out the BIV hypothesis *before* we are entitled to any knowledge claims, as, *ex hypothesi*, there can be no grounds – bar its ‘logical possibility’ – for believing that the sceptical hypothesis might be true.

VI Conclusion

If what I have argued in this paper is correct, no solution need be found to ‘closure-based’ sceptical arguments, since it turns out that there can be no such thing. Given that we have seen that the closure principle itself does no real sceptical work, as it only comes into play *after* one has already conceded knowledge (or concluded that this isn’t possible), everything hangs on our reasons for accepting (BIV2) – reasons that we have seen to depend, on the ‘default view’, on a prior endorsement of IA and RIT. If we accept this argument, then we get the conception that all we can ever have is (non-factive) cookie-experiences, in which case we don’t need to deny closure, since we cannot, in any case, deduce from those that we are not confronted by ‘vat-cookies’.

So, it seems that if one wants to make any headway with radical scepticism, the way to do this is not to attack an innocent principle, but rather to subject to closer scrutiny the ‘default view’ of perceptual reasons, since it is this conception, not the closure principle, that provides ammunition for the sceptical thought that our perceptual experiences *as a class* always fall short of providing one with knowledge of an independent physical world. That is to say, we need to look more closely at the reasons one might have for thinking, like Dretske, that one is disbarred, *ab initio*, from ever perceiving real (physical) cookies. For it is this notion that has the sceptical implications – if we start with this assumption, we get scepticism for free; no closure-based sceptical argument is necessary. It is not surprising, therefore, that Dretske’s apparently new-fangled case against closure collapses, in the end, into an old-fashioned form of veil-of-appearances scepticism.

References

- Atkins, P., and Nance, I. (2014), 'A Problem for the Closure Argument', *International Journal for the Study of Scepticism* 4, 36-49.
- Bennett, J. (1971), *Locke, Berkeley, Hume: Central Themes* (Oxford: Oxford University Press).
- Brueckner, A. (1994), 'The Structure of the Skeptical Argument', *Philosophy and Phenomenological Research*, 54, 827-35.
- Burge, T. (2003), 'Perceptual Entitlement', *Philosophy and Phenomenological Research*, 67, 503-48.
- Brueckner, A. (2010), 'Skepticism and Closure' in Jonathan Dancy, Ernest Sosa, and Matthias Steup (eds), *A Companion to Epistemology* (Oxford: Blackwell), 3-12.
- Coliva, A. (2012), 'Moore's Proof, Liberals and Conservatives – Is There a (Wittgensteinian) Third Way' in Coliva (ed.), *Mind, Meaning and Knowledge* (Oxford: Oxford University Press), 323-351.
- Conee, E. (2007), 'Disjunctivism and Anti-Skepticism', *Philosophical Issues* 17, 16-36.
- Dretske, F. (1970), 'Epistemic Operators', *Journal of Philosophy* 60, 1007-23.
- Dretske, F. (1971), 'Conclusive Reasons', *Australasian Journal of Philosophy* 49, 1-22.
- Dretske, F. (2005a), 'The Case against Closure' in Ernest Sosa and Matthias Steup (eds.), *Contemporary Debates in Epistemology* (Oxford: Blackwell).
- Dretske, F. (2005b), 'Reply to Hawthorne' in Ernest Sosa and Matthias Steup (eds.), *op.cit.*

- Dretske, F. (2010), 'Epistemological Self-Portrait' in Jonathan Dancy, Ernest Sosa, and Matthias Steup (eds), *A Companion to Epistemology* (Oxford: Blackwell).
- Hawthorne, J. (2005), 'The Case for Closure' in Ernest Sosa and Mathias Steup (eds.), *op.cit.*
- Klein, P. (2004), 'Closure Matters: Academic Scepticism and Easy Knowledge', *Philosophical Issues* (Epistemology) 14, 165-84.
- McGinn, C. (1984), 'The Concept of Knowledge', *Midwest Studies in Philosophy* IX, 529-544.
- Millar, A. (1991), *Reasons and Experience* (Oxford: Clarendon Press).
- Millar, A. (2011), 'Disjunctivism and Scepticism' in J. Greco (ed.), *The Oxford Handbook of Scepticism* (Oxford: Oxford University Press).
- Nagel, T. (1986), *The View from Nowhere* (Oxford: Oxford University Press).
- Nozick, R. (1981), *Philosophical Explanations* (Oxford: Oxford University Press).
- Pollock, J. (1974), *Knowledge and Justification* (Princeton: Princeton University Press).
- Pritchard, D. (2005), 'The Structure of Sceptical Arguments', *Philosophical Quarterly*, 55, 37-52.
- Pritchard, D. (2007), *Epistemic Luck* (Oxford: Oxford University Press).
- Pritchard, D. (2008), 'McDowellian Neo-Mooreanism' in A. Haddock and F. Macpherson (eds), *Disjunctivism: Perception, Action, Knowledge* (Oxford: Oxford University Press), 283-310.
- Pritchard, D. (2009), 'Wright contra McDowell on Perceptual Knowledge and Scepticism', *Synthese* 171: 467-79.

- Pritchard, D. (2012), *Epistemological Disjunctivism* (Oxford: Oxford University Press).
- Stroud, B. (1984), *The Significance of Philosophical Scepticism* (Oxford: Oxford University Press).
- White, R. (2014), 'What is my Evidence that here is a Hand?' in D. Dodd and E. Zardini (eds.), *Scepticism and Perceptual Justification* (Oxford: Oxford University Press), 298-321.
- Williams, M. (1996), *Unnatural Doubts* (Princeton: Princeton University Press).
- Wright, C. (2002), 'Anti-Sceptics Simple and Subtle: G. E. Moore and John McDowell', *Philosophy and Phenomenological Research* 65, 330-48.
- Wright, C. (2014), 'On Epistemic Entitlement (II): Welfare State Epistemology', in D. Dodd and E. Zardini (eds.), *op.cit.* 213-247.
- Yalcin, Ü. (1992), 'Sceptical Arguments from Underdetermination', *Philosophical Studies* 68, 1-34.