

No Spindly Brown Grass: Knowledge, Closure and Subjective Distinguishability

Genia Schönbaumsfeld

(Forthcoming in Matthew Jope and Duncan Pritchard (eds.), *New Perspectives on Epistemic Closure* (London: Routledge).)

I Introduction

Conceptions of perceptual reasons in contemporary epistemology tend to be scepticism-driven. That is to say, epistemologists implicitly (or explicitly) take the sceptical problem as a point of departure and then formulate a notion of perceptual grounds that, right from the start, makes serious concessions to scepticism¹. This constitutes, in effect, a ‘lowest common denominator’ approach – the resulting view is whatever one might be left with after the sceptic (or our sceptical *alter ego*) has done her work. But if that is the strategy, then it is not surprising that the outcome, to put it in Dretske’s memorable words, is that ‘philosophy is a business where one learns to live with spindly brown grass in one’s own yard because neighbouring yards are in even worse shape’ (Dretske 2005: 43).

In this paper, I want to argue that to take such an approach is a mistake. We should rather proceed the opposite way. In other words, we should begin with a conception of perceptual reasons that fits the facts and our ordinary intuitions, and then explore what the

¹ This seems true of fallibilism, contextualism, contrastivism, and accounts that deny closure. Without the sceptical problem as a motivating background force, many of these conceptions would lose their appeal (i.e. would not strike us as independently plausible).

ramifications for scepticism might be. We should not, *ab initio*, cede the high-ground to the sceptic and let her dictate the terms of engagement.

I will proceed as follows. In order to show the scepticism-driven nature of much contemporary epistemology, I will make use of Dretske's attempt to deny closure as a case study. For although most epistemologists have not been impressed by Dretske's manoeuvre, the reasons that drive him to make it, are widely shared. So, if I am right about the underlying diagnosis, the lesson will generalize and enable us to resist certain apparently compelling moves that in fact drive us right into the arms of scepticism. This is not yet to construct a full anti-sceptical strategy, but merely an attempt to undermine ways of thinking that are less compulsory than they might seem, and more scepticism-friendly than strictly required.

II Scepticism and Closure

In his paper, 'The Case Against Closure', Dretske makes the following claims:

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)).

This is quite a striking passage. Is it as self-evident as Dretske would like us to think that one ‘surely isn’t claiming to see that there are physical objects in claiming to see there are cookies in the jar’? For, *pace* Dretske, one’s ordinary way of thinking about cookies does precisely seem to be that they are sweet, mind-independent physical objects. Imagine that you ask your friend whether she believes that there are cookies in the jar (when there are), and she says ‘yes’. Further imagine you then go on to ask her, ‘do you believe that these cookies exist only in your mind?’, and I’m confident, she would say ‘no’. Imagine further still that you then ask her whether she thinks that seeing that there are cookies in the jar is compatible with believing that she is hallucinating that there are, and I’m sure she would also say ‘no’. Of course, one doesn’t settle philosophical questions by asking one’s friends. Nevertheless, what this exchange shows is that Dretske’s proposal appears quite revisionary – he is denying we would want to claim what most normal people, and, indeed, most philosophers (when not mesmerized by scepticism) would want to say; namely, that to see that there are cookies in the jar is to see certain kinds of physical objects in the jar.

So, why does Dretske want to say something that (certainly at first glance, and perhaps later too) seems highly counter-intuitive? The reason is that he is already convinced by a certain sceptical argument. After all, he says, ‘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.’ In other words, Dretske appears to be endorsing the conclusion of the ‘Indistinguishability Argument’, which runs as follows:

The Indistinguishability Argument

Q1 In the bad case it only seems (looks) to me as if P.

Q2 I cannot subjectively distinguish P (the good case) from the mere appearance as of P (the bad case).

C1 If I cannot subjectively distinguish between P and the mere appearance as of P, I cannot know that P (that I am in the good case).

C2 I cannot know that P.

That is to say, because I cannot visually distinguish a real cookie from the hallucination of one – they both ‘look’ the same – I cannot know that the ‘cookie’ in the jar in front of me is a real (i.e. physical) cookie.

This sounds like a sceptical conclusion. Interestingly enough, however, Dretske wants to persuade us that it is not, and this is where his (in)famous denying of the closure principle² comes in. Dretske reasons as follows. When we perceive that P, P has certain ‘heavyweight’ implications – such as that cookies are mind-independent physical objects – 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 (Dretske 2005a: 16)³. So, 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

² The closure principle states that ‘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’ (I am here following Hawthorne (2005) and Pritchard (2012)).

³ Or, to put it in 1970’s speak: the epistemic operator ‘to know’ does not penetrate to all the known logical consequences of what is known (see Dretske (1970)).

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 (Dretske 2005a: 17): ‘The only way to preserve knowledge of homely truths, the truths everyone takes themselves to know, is...to abandon closure’ (Dretske 2005a: 18)⁴. In other words, denying closure is a last-ditch attempt to ward off the scepticism that, according to Dretske, would otherwise be inescapable.

Let’s flesh out Dretske’s account a bit more. In his paper, ‘Epistemic Operators’, where Dretske first developed these ideas, he explains that there are certain presuppositions associated with a statement, which, ‘although their truth is entailed by the truth of the statement, are not part of what is operated on when we operate on the statement with one of our epistemic operators. The epistemic operators do not *penetrate* to these presuppositions’ (Dretske 1970: 1014)⁵. For example, I go to the zoo and look at an enclosure labelled ‘zebras’. I see animals that look like zebras. I conclude that there are zebras in the pen. However, can I know that these ‘zebras’ are not mules cleverly disguised to look like zebras? No. That they are not mules cleverly disguised to look like zebras is a presupposition that I’m ruling out because it is highly unlikely. But can I know that this presupposition is true? No. My perceptual evidence, according to Dretske, is not sufficient to rule out this possibility (Dretske 1970: 1016).

⁴ 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.’

⁵ Crispin Wright (2002, 2014) has a similar conception: unless we presuppose that there is an external world, we cannot know any perceptual claim to be true. Wright puts this in terms of his I-II-III schema:

I. My current experience is in all respects as if P.

II. P

III. There is an external world.

With this schema in place, one can only know (II) if one already knows (III). See *The Illusion of Doubt*, chapter 2, for further discussion.

Again, this sounds like a sceptical conclusion. So, how are we supposed to avoid scepticism? Dretske agrees with the sceptic that you cannot know that the zebras are not cleverly disguised mules. But he believes that this does not imply that we do not know that the animals are zebras, because he rejects the principle that gets us to this conclusion: the principle that if you do not know that Q is true, when it is known that P entails Q, then you do not know that P is true (Dretske 1970: 1016). In other words, Dretske rejects the principle that knowledge is ‘closed’ under known logical entailment.

This, to put it mildly, generates some very odd consequences. It implies, for instance, that I can know that the zebra in the pen is not a mule, as my perceptual evidence – its looking like a zebra – speaks against the relevant alternative that it’s a mule. Nevertheless, I cannot know that the zebra is not a *cleverly disguised mule*, as my perceptual evidence does not speak against this scenario. Hence, I can know that the animal in the enclosure is not a mule, but I cannot know that it is not a cleverly disguised mule. And this seems not just to be an ‘abominable conjunction’⁶, but courts downright logical contradiction: I can both know and not know that the animal in the enclosure is a mule. Similarly, I can know, on Dretske’s conception, that the wall is red and not white, but I cannot know that the wall is not white cleverly illuminated to look red. This sounds like a case of wanting to have one’s logical cake and eat it too.

One might try to alleviate the tension by confining oneself to ‘looks’ talk. That is to say, if all that Dretske means by saying we know the animal is a zebra, not a mule, is that we know that the animal *looks* like a zebra, then there is no logical contradiction. For that the animal *looks* like a zebra is, of course, perfectly compatible with the animal in fact being a cleverly disguised mule. Unfortunately, however, if we take this route, it seems that we are

⁶ Conjunctions of the sort ‘I know I have a hand, but I don’t know I’m not a handless BIV’ (see DeRose (1995)). DeRose makes this point against Nozick.

not left with much knowledge at all. For how can something's *looking* like a zebra be sufficient for knowing it's a zebra? All sorts of things might *look* like a zebra and not be a zebra. My child's toy zebra looks like a zebra and is not a zebra⁷. A tiger might look like a zebra if my vision is obscured by branches and sufficiently far away etc. Consequently, Dretske seems to be faced with the following dilemma: either his account is logically problematic (I can both know and not know that an animal is not a mule), or we can't know much more than how an object happens to look. And a sceptic would, naturally, be quite happy with either claim.

Does Dretske's later, 'information-caused belief' conception (Dretske 1981) save him from being impaled on this dilemma? It appears not. For even if we were to agree that the animal's looking like a zebra causes one to believe that the animal is a zebra and not a mule, it is still going to be the case that this information underdetermines whether the animal is a mule cleverly disguised to look like a zebra. Consequently, on this conception too, I'm going to come out as embracing both that I can know that the animal in the pen is a zebra and not a mule, but I cannot know that the animal is not a cleverly disguised mule. Hence, in order to avoid this problem, Dretske would again have to retreat to 'looks' talk – all one can know is that the animal looks like a zebra; nothing more, nothing less. But, if so, he cannot really avoid the sceptical conclusion.

If, strictly speaking, all we can know, on Dretske's conception, is that it looks like there are cookies in the jar or that there are zebras in the pen, then the information that we are receiving from our perceptual experiences always underdetermines whether these experiences are of actual physical cookies or merely figments of our imagination. In other words, our perceptual experiences, on this view, will always be experiences of something which leaves it

⁷ This is a point that Dretske should concede, given that he says that a decoy duck is not a duck (Dretske 1981: 57)!

open whether these experiences are of actually existing physical objects or not. And this implies that whether we find ourselves in the good or the bad case, our perceptual reasons are the same⁸; they always underdetermine whether we are confronted by an actual physical object or are merely hallucinating it. In *The Illusion of Doubt*, I called this view, which is shared by the vast majority of contemporary epistemologists, the Reasons Identity Thesis (RIT)⁹. Endorsement of this thesis has the following unwelcome implications for Dretske's case against closure. Consider a standard closure-based sceptical argument:

Standard Closure-Based Sceptical Argument

(BIV1) If I know I have two hands (or that there is a cookie or zebra in front of me), 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.'). But, it seems,

(BIV2) I don't know I'm not a BIV.

⁸ I take this to be an epistemological conclusion, not a metaphysical one, although it does have metaphysical implications – namely the essentially inferential nature of perceptual knowledge. For more on this, see *The Illusion of Doubt*, chapter 3.

⁹ See, for example, Avnir (2019), Burge (2003), Conee (2007), Coliva (2012, 2014, 2015), McGinn (1984), Millar (1991), Nagel (1986), Pollock (1974), Pryor (2000), Stroud (1984), White (2014), Wright (2002, 2014).

So,

(BIV3) I don't know I have two hands (or that there is a cookie or zebra in front of me).

Dretske wants to avoid (BIV3) by denying the closure principle. That is to say, he contends that he can know that he has hands (or that there are cookies in the jar) without this knowledge transmitting to whatever is logically entailed by this knowledge. But, as we have already seen, this in turn implies that Dretske must confine himself to knowledge of how things appear (how they look), otherwise, he could neither avoid the logical contradictions discussed above, nor make plausible why knowledge of the presence of cookies does not confer knowledge of sweet physical objects (since everyone agrees that cookies *are* sweet physical objects).

If so, however, then the antecedent in (BIV1), on Dretske's parsing, doesn't mean what it ordinarily means. For, on Dretske's reading, 'knowing one's got hands (cookies etc.)' means 'knowing that one has an experience as of a hand' (knowing that there is something that looks like a hand), not 'knowing that there is an actual physical hand there' (since this, on Dretske's view, is a 'heavyweight implication' that we cannot know). But, if so, (BIV1) needs to be modified to reflect this fact. In other words, (BIV1) should read:

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

If this is what (BIV1) really means, however, then it generates the unfortunate consequence that the premise now turns out to be false: one cannot deduce that one is not a BIV from the presence of an experience as of there being a hand in front of one. Experiences as of hands are non-factive and don't imply the presence of real, physical hands¹⁰. If (BIV1)* is false, though, it naturally cannot serve as a premise in a closure-based sceptical argument. Hence, once this is made clear, it becomes apparent that, on Dretske's conception, no closure-based sceptical argument can run and, therefore, that a denial of closure is entirely redundant¹¹. Given that Dretske is already *starting* with the claim that one can only have knowledge of hand-like appearances (experiences as of there being hands), never of the (physical) hands themselves, and one cannot deduce something from a proposition that isn't entailed by it, the denying 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 fact that one is seeing a physical hand, as only knowledge of a physical hand entails the falsity of the sceptical conclusion. But that is not what Dretske is denying. Dretske is denying that one can deduce that one is not a BIV from a hand-like appearance, and this is both true and entirely consistent with the closure principle.

In other words, the reason why I cannot, for example, reason from 'I see that there is a cookie in the jar' to the presence of a physical 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. because he has already accepted the Reasons Identity Thesis). 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 not sufficient to ground knowledge of 'heavyweight' cookies (about the fact that there are real cookies there).

¹⁰ Also see Wright (2002, 2014).

¹¹ I first argued this in Chapter 1 of *The Illusion of Doubt*.

In his epistemological self-profile (Dretske 2010), Dretske comes very close to admitting this. He says: ‘If I can’t simply assume that deceptive demons and cunning hoaxes are not possibilities my evidence need rule out, possibilities I can ignore...then I really can’t see what I purport to see – that my wife is on the sofa. I can no longer see, just by looking, that she is there because these ways of her not being there would look exactly the same way to me...So closure fails in these special (isolated) cases because the key premise, P, is not available to those who want to use it to infer Q’ (Dretske 2010: 134). The latter sentence is revealing. If P is not available to those who want to use it to infer Q, then it is not that closure fails. Rather, the fact that I don’t know that P means that I can’t use knowledge of P to infer that Q. And this is exactly what one would expect; exactly what the closure principle says: If I don’t know that Q, this is because I don’t know that P.

Given that all I can know, on Dretske’s view, is that I have an experience as of my wife’s being on the sofa, and this is compatible with being deceived about this experience, it follows that Dretske, as he himself admits, cannot know that his wife is really on the sofa (it only looks/seems/appears that way). Consequently, the blame for the sceptical implication lies squarely with Dretske’s antecedent endorsement of the Reasons Identity Thesis, not with the closure principle, which only comes into play if we have already established that knowledge of physical objects is possible. Since this is deemed impossible by Dretske *ab initio*, however, no knowledge of P is, in fact, available from which one could infer Q. For all one can know, on Dretske’s conception, is that it appears to one that P – and from an appearance that P it is naturally not possible to deduce that Q.

The closure principle, therefore, cannot furnish one with new empirical knowledge, but only informs one of what follows logically from what one already knows. Hence, 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. Conversely, if one's putative knowledge that P is confined to knowledge of its merely seeming to one that P, then one cannot, in any case, derive an anti-sceptical conclusion from such an appearance, and so a denial of closure is obsolete.

So, instead of scapegoating an innocent principle, what we need to do is to examine the reasons why Dretske (and others) believe that we can never know more than how things look (seem). That is to say, we need to subject the Indistinguishability Argument and the conception of perceptual reasons that it appears to thrust upon us to more detailed scrutiny, as it is acceptance of these notions, not of the closure principle, that pushes us towards scepticism.

III Experiential Markers and Subjective Distinguishability

As we have seen, Dretske's master argument consists in claiming that by seeing cookies in a jar, we can never come to know that there really are sweet physical objects in the jar, as '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' (Dretske 2005: 14). Since actual cookies and hallucinations of cookies are phenomenologically indiscriminable, in other words, Dretske thinks that we can never know that we are confronted by real cookies rather than mental figments.

Whatever else is wrong with this reasoning (of which more below), the first thing to note here is that this conclusion is a bit quick. It hardly goes without saying that if I cannot tell the difference by 'looking' – where I presume this means 'introspectively distinguishing

on the basis of phenomenology alone’ – I can never know that what I am looking at is real and not a hallucination. For even if I cannot tell the difference on the basis of how things appear to me, I could conduct further empirical investigations (such as, say, checking for paint in the zebra case), or looking more closely; or consulting other, maybe better trained, people. So, perhaps Dretske is assuming that these routes are not available to me, because they are no more ‘secure’ than looking is – where ‘secure’ probably also means that I already need to have a way of phenomenologically distinguishing *those* cases from their ‘fake’ counter-parts. But this seems question-begging, if one is not already convinced of the conclusion of the Indistinguishability Argument.

Second, it is entirely unclear whether what Dretske (and the vast majority of epistemologists) thinks we need – subjective phenomenological discriminability – would in fact help us, even if it were available. For if it wouldn’t, then we should stop thinking that not having subjective discriminability is, by itself, sufficient to make the Reasons Identity Thesis and the conclusion of the Indistinguishability Argument compulsory¹².

So, let’s suppose, for the sake of argument, that an experiential marker were available that would enable us subjectively to distinguish a real cookie from a deception. To this end, let’s imagine that non-veridical perceptual experiences always come along with a pinkish tinge, so that whenever one appears to perceive something that is out of kilter with the way the world is, one’s experiences have this hue. For example, when I seem to see a cookie in a jar which is actually a hologram, I will experience the scene in front of me as if I were looking at it through rose-tinted glasses. This would alert me to the fact that the hologram only looks like a cookie, but is not in fact a cookie. Would this be a useful result? Contrary to what one might expect, it seems not, for the following reasons.

¹² For more on this, see *The Illusion of Doubt*, chapter 1, as well as Pritchard (2012).

If we want the marker in question to alert us to every occasion of something's appearing other than how it actually is, then virtually all our perceptual experiences will end up coming with a rose-coloured hue. One might call this the problem of 'pinkish encroachment'. The reason for this is the fact that we have a unique point of view on the world and every time we change our vantage point, things will look different to us. Large towers look small every time we move away from them; square towers look round the greater the distance from which we view them. In order for everything not to appear pink to us for most of the time, therefore, the rose-tinted marker would, at the very least, have to be able to track how things seem to us under standard conditions, and only light up when we are straying from those. But even with such a mechanism built in, we would end up having a lot of rose-tinted experiences. For example, the marker would have to be 'on' whenever we looked at something further away in our visual field, or in motion and so on. Furthermore, it would have to be able to 'guess' which object in our environment we are focussing on, in order not to light up all the time, given that some objects in our visual field will look the way they are, whereas others, because they are further away, will not. Since such an 'intentional marker' seems impossible, it is hard to see how any experiential marker could be fine-grained enough to convey useful information¹³.

What is more, without a distinction between how things look and how they really are – the apparent 'gap' exploited by the sceptically-minded that the marker is supposed to close – a creature with a point of view on a world could not learn anything about this world. An object, for instance, *ought* to look small from afar – if it stayed large even though we were moving away from it, something would be seriously wrong. For it would imply that the object is growing as we are moving away, that our eye-sight is failing, that we are

¹³ If, on the other hand, one tried to 'correct' the marker for every change brought on by changing perspective (e.g. by moving), then it would virtually never be on. In other words, it seems that experiential markers would either have to be on all the time or never, and neither option is compatible with information transmission.

hallucinating, that distance is having an uncanny effect on objects, that my point of view in space can affect the nature of an object (not just its appearance) etc. In other words, there are ways an object ought to look from a given perspective, despite that look's not necessarily being indicative of what the object is really like. Since this is an ordinary feature of perception, it would be useless, and a case of more pinkish encroachment, if the experiential marker came on every time things looked just how they ought to look from one's perspective.

In order to get round these problems, it might, perhaps, seem a good idea to distinguish between 'local' and 'global' experiential markers. 'Local' experiential markers that are supposed to inform us of deceptive appearances in our immediate environment don't, as we have just seen, appear capable of serving the purpose they were originally drafted in for, as, due to pinkish encroachment, they are incapable of conveying much useful information. The marker might, on occasion, be able to pick up on a genuine perceptual illusion – e.g. a hallucination induced by a drug – but since this marker would have come on many times before when nothing was amiss except that we had changed position and perspective, we would not be able to make much use of this information (for we wouldn't be able to distinguish this case from all the other cases where the marker had been present). But perhaps a 'global' experiential marker that came on only if we were in radical or global sceptical scenarios, would be more promising.

Imagine, for example, that in a BIV-world all of our apparent perceptual experiences come with a pinkish hue, since none of them are veridical. Would this in fact help the BIV? Only if the BIV could come to know that experiential rose-tintedness indicates non-veridicality, and it is hard to see how the BIV could come to know this. For if the BIV were a BIV from birth and all its perceptual experiences were rose-tinted from the off, then it would regard rose-tintedness as just a normal feature of perceptual experience (as it would never have had any non-pink perceptual experiences). If, on the other hand, the BIV had previously

been a normal person and woke up to find itself having rose-coloured perceptual experiences, then it would probably think something was wrong with its perceptual faculties, not entertain the idea that it had surreptitiously been envatted. Conversely, if the BIV suddenly became disenvatted and started having non-pinkish perceptual experiences, then it would similarly think that something had happened to its perceptual capacities, not that it had been released, all of a sudden, into the real world. So, whichever way we look at it – whether from a ‘local’ or a ‘global’ perspective – experiential markers don’t seem able to convey much information¹⁴.

If this is right, then it looks like the presence of experiential markers – the only conceivable way of fleshing out the notion of subjective introspective distinguishability (discriminability on the basis of phenomenal content alone) – would not advance our case against scepticism, since these markers (if such could be had) would not be able to ‘tell’ us what we wanted them to. Consequently, we need to be very cautious about agreeing that we cannot know that the cookie in the jar is a real cookie, *unless* such markers (and subjective introspective distinguishability) were available, since it now turns out that such markers would be entirely useless to the task at hand¹⁵.

IV Implications for Dretske’s Argument and an Alternative Conception

¹⁴ For more on this, see Schönbaumsfeld (forthcoming).

¹⁵ Now perhaps some would regard this as a way of compounding scepticism: If experiential markers cannot give one perceptual knowledge, nothing can. But that would be irrational. If experiential markers are not a means to the end they were envisaged for, then one ought to consider an alternative means to this end, rather than come to the conclusion that they are the only – albeit logically impossible – game in town. Since an alternative conception is available (see the next section), there is no reason to throw in the towel at this point.

We have seen 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 jar), 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)). This may well be true, but the question is why one should accept the view that your reasons for believing P – say, seeing that there are cookies in the jar – are reasons that you would have even when P is false. For, given the foregoing considerations, this is not as obvious as Dretske seems to think, since one might believe instead that in cases where P is false, it would only *appear as if* one had the reasons to believe that P that one has in the good case, not that, *in actual fact*, one has those reasons (whether this misleading appearance means that one lacks any reason at all to believe that P is a further question, but not one that needs to be adjudicated here). 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.

Why does Dretske not consider this possibility? Again, because he is implicitly relying on the 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 hologram, I cannot know that I am experiencing an actual (physical) cookie. But, as we have already seen, this conclusion doesn't follow unless one is already convinced of the Reasons Identity Thesis. Since one key argument for accepting this thesis – the absence, in perception, of experiential markers – has just collapsed (see previous section), there is nothing that forces us to adopt this conception.

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. For good further discussion of this point, see Pritchard (2012: 42-5).

Consequently, it remains a live option to think that because the good case is better than the bad case, it also furnishes one with much better reasons than its counterpart: in the good case, one's cookie-experiences can be factive (i.e. entail the presence of an actual physical object), while, in the bad case, this is not so. To perceive that there is a cookie in the jar is not, therefore, to perceive something which leaves it open whether I am perceiving a cookie or am merely hallucinating it. For if I'm 'perceiving' a 'mental cookie', I'm not, of course, *perceiving* anything (and certainly not a 'type of cookie'¹⁷). At best, I'm being appeared to in a cookie-like manner. To think otherwise, is just to adopt what McDowell (1998a, 1998b) calls the Highest Common Factor conception and I call the Reasons Identity Thesis: the notion that all that perception can give us is appearance-based reasons that are compatible with those appearances being massively deceptive.

Since there is no good reason to accept this conception, however, we also don't need to endorse Dretske's claim that we can *never* 'perceive heavyweight implications to be so'. Rather, we could 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 in the jar, then I could also come to know – in virtue of my factive cookie-experience – that I am not seeing a 'hologram-cookie' in the jar, as seeing a 'hologram-cookie' is to be taken in by a simulation, not to perceive (see) a sweet physical object¹⁸. But coming to know that I am not seeing a 'hologram-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 of a physical cookie, and not, *pace* Dretske, an experience of

¹⁷ For 'hallucinated cookies' are not tokens of the type 'cookie'. Dretske sometimes speaks as if they were. But if he literally meant that, the view would be absurd.

¹⁸ 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.

something that leaves it open whether it is the experience of a physical cookie or of a ‘hologram-cookie’ (since, on such a conception, no cookie-experience could ever be factive – that is to say, entail the presence of an actual physical cookie, not the mere appearance of one).

Of course, Dretske is right that one cannot ‘perceive’ that one is not perceiving a ‘hologram-cookie’, since all one can perceive is a *cookie*. But unless one has already ruled in advance that ‘cookie-perception’ can *never* give one knowledge of physical cookies, perceiving a cookie can, in the good case, be sufficient to give one knowledge of the presence of physical cookies even if one can’t subjectively distinguish between physical cookies and ‘hologram-cookies’¹⁹. Consequently, one doesn’t, *pace* Dretske, need an *additional* reason to believe the ‘heavyweight’ implication that one isn’t perceiving a ‘hologram-cookie’. In the good case, I can know non-inferentially that there is a cookie before me; I do not need to infer its existence from experiences that are neutral between being either experiences of actual cookies or ‘hologram-cookies’.

What is more, Dretske’s argument to the contrary seems to be relying on the principle that if one perceives that P, and one knows that P entails Q, one cannot know that Q if Q is not itself something one can perceive. But that principle is false. If I know that Pierre murdered Maria, because I saw him kill her, I also know that Felipe, a friend of Maria’s, was not the murderer, although I cannot (do not) ‘perceive’ that Felipe was not the murderer (since I have no perceptions about Felipe). Hence, although knowledge of Felipe’s not being the murderer is entailed by my having perceived that Pierre is the murderer, my knowledge of Felipe’s not being the murderer does not depend on my having ‘perceived’ that Felipe is not the murderer. Similarly, in order to know that I’m seeing a real cookie when I am, I don’t

¹⁹ Compare Klein (2004), who makes a similar point about zebras and cleverly disguised mules; see also Pritchard (2012), parts II and III.

need to be able to do two different things: to perceive a cookie and, on top of that, to perceive that a cookie is a physical object (whatever that could so much as mean).

So, if I am in the good case, no further evidence is required to license the move from ‘I see that there is a cookie in the jar’ to ‘I see that there is a physical 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). Therefore, *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).

Finally, were we to grant that seeing that there is a cookie in the jar is *never* a reason to believe that we are seeing actual physical cookies in the jar (since this is an ‘unperceivable’, ‘heavyweight’ implication), then it’s unclear why we should nevertheless regard this as a good reason to believe that there are *cookies* in the jar (as opposed to anything, or nothing). An appeal to ‘relevant alternatives’ is not going to head off this objection, as one doesn’t know what something is just because one can distinguish it from salient alternatives²⁰.

At this point, one might try a last defence measure: to construe Dretske’s view as an early articulation of present-day contrastivism (Schaffer 2005). In other words, perhaps we should regard Dretske as saying that knowledge is ‘contrastive’ – that is, that there really is no such thing as knowing that there is a cookie *simpliciter*, for example, as knowing that P is in fact a three-place relation concerning a subject, a proposition, and a contrast class. Consequently, one can only know that there is a cookie in the jar relative to the relevant

²⁰ Let’s say I come across an object that I haven’t encountered before, which looks like a tool or piece of equipment. I clearly won’t find out what it is just because I know it’s not one of the relevant alternatives: a hammer, a chain-saw, a pair of pliers etc.

contrast class – say, that there is a cookie-rather-than-a-pretzel in the jar – and not that there is a cookie in the jar *tout court*.

But to adopt such a strategy reaps few rewards. For, on the contrastivist's proposal too, one will come out as not knowing very much. Hence, there doesn't seem to be much payoff for espousing such a revisionary stance about knowledge.

To see this, let's take a closer look at the contrastivist's conception. According to Schaffer, contrastivism enables one to know that one has hands rather than stumps, but not to know that one has hands rather than vat-hands (Schaffer 2005: 262): 'Since the existence of possibilities outside one's discriminatory range does not imply the absence of any possibilities inside that range, sceptical doubts do not imply the absence of any ordinary knowledge' (Schaffer 2005: 263). Do they not, though? For what are we to make of this 'ordinary knowledge'? As we already saw at the beginning, 'ordinary' people believe that ordinary cookies are physical cookies, and not simulations or hallucinations. Does Schaffer's account take this on board any more than Dretske's does? No.

For if I cannot know that the cookie before me is a cookie and not a vat-cookie, as this lies outside my discriminatory range, then neither can I know that the cookie before me is a real (i.e. physical) cookie, as opposed to the mere appearance of one. Consequently, just as on Dretske's conception, whatever 'ordinary knowledge' I can possess (knowledge of 'cookies-rather-than-pretzels'; 'hands-rather-than-stumps') will at best be knowledge of appearances only (i.e. knowledge that never rules out the contrast class of deceptions or hallucinations). But, if so, we are back to square one and the sceptic (or our sceptical *alter ego*) has the last laugh.

V Conclusion

In summary, in this paper I have argued for four main claims:

- 1) Acceptance of the Reasons Identity Thesis, which is widely endorsed in contemporary epistemology, is already to make substantial concessions to scepticism that are unwarranted.
- 2) Implicit endorsement of this thesis leads Dretske to deny closure – a move that is a red herring and that won't, in fact, preserve us from scepticism.
- 3) The absence of subjective introspective discriminability is not a good reason to adopt the Reasons Identity Thesis, as experiential markers, if, *per impossibile*, they were available, is not a means to avoiding scepticism.
- 4) An alternative, less revisionary, conception of perceptual reasons is available that avoids the problems generated by RIT and that does not sell out to scepticism from the outset.

Of course, what I have done in this paper constitutes only a first step towards a full anti-sceptical strategy²¹. Nevertheless, I hope to have persuaded the reader that we do not, right from the start, need to concede that we are 'walled off' from the facts and that knowledge of 'heavyweight' propositions is in principle impossible. If this is right, we don't need to settle for the 'spindly brown grass' that Dretske speaks of: without the Reasons Identity Thesis

²¹ See *The Illusion of Doubt* for the full picture.

throwing a spanner into the works from the very beginning, verdant regeneration can beckon still²².

References

- Avnur, Y. (2019), 'The Nature and Limits of Skeptical Criticism', *International Journal for the Study of Skepticism* 9/3, 183-205.
- Burge, T. (2003), 'Perceptual Entitlement', *Philosophy and Phenomenological Research*, 67, 503-48.
- 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.
- Coliva, A. (2014), 'Moderatism, Transmission Failures, Closure and Humean Scepticism' in Dodd and Zardini (eds), *op.cit.*, 248-76.
- Coliva, A. (2015), *Extended Rationality: A Hinge Epistemology* (London: Palgrave Macmillan).
- Conee, E. (2007), 'Disjunctivism and Anti-Skepticism', *Philosophical Issues* 17, 16-36.
- DeRose, K. (1995), 'Solving the Sceptical Problem', *Philosophical Review* 104: 1, 1-52.
- Dretske, F. (1970), 'Epistemic Operators', *Journal of Philosophy* 60, 1007-23.
- Dretske, F. (1981), *Knowledge and the Flow of Information* (Boston: MIT Press).

²² I would like to thank Duncan Pritchard for helpful comments on an earlier draft.

- Dretske, F. (2005a), 'The Case against Closure' in Ernest Sosa and Matthias Steup (eds.), *Contemporary Debates in Epistemology* (Oxford: Blackwell), 13-25.
- Dretske, F. (2005b), 'Reply to Hawthorne' in Ernest Sosa and Matthias Steup (eds.), *op.cit.*, 43-6.
- Dretske, F. (2010), 'Epistemological Self-Profile' 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.* 26-42.
- McDowell, J. (1998a), 'Criteria, Defeasibility and Knowledge' in McDowell, J., *Meaning, Knowledge and Reality* (Cambridge, MA: Harvard University Press), 369-94.
- McDowell, J. (1998b), 'Knowledge and the Internal' in *Meaning, Knowledge and Reality*, 395-413.
- McGinn, C. (1984), 'The Concept of Knowledge', *Midwest Studies in Philosophy IX*, 529-544.
- Millar, A. (1991), *Reasons and Experience* (Oxford: Clarendon 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. (2012), *Epistemological Disjunctivism* (Oxford: Oxford University Press).
- Pritchard, D. (2015), *Epistemic Angst* (Princeton: Princeton University Press).
- Pryor, J. (2000), 'The Skeptic and the Dogmatist', *Noûs*, 43, 517-49.

Schaffer, J. (2005), 'Contrastive Knowledge' in Tamar Szabo Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology volume 1* (Oxford: Oxford University Press).

Schönbaumsfeld, G. (2016), *The Illusion of Doubt* (Oxford: Oxford University Press).

Schönbaumsfeld, G. (forthcoming), 'Subjective Distinguishability – Who Needs it?' in Yuval Avnur (ed.), *Midwest Studies in Philosophy*.

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.