THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Relevant Alternatives, Perceptual Knowledge and Discrimination

Citation for published version:

Pritchard, D 2010, 'Relevant Alternatives, Perceptual Knowledge and Discrimination', *Noûs*, vol. 44, no. 2, pp. 245-268. <<http://onlinelibrary.wiley.com/doi/10.1111/j.1468-0068.2010.00739.x/abstract>>

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Noûs

Publisher Rights Statement:

© Duncan Pritchard, 2010. Pritchard, D. (2010). Relevant Alternatives, Perceptual Knowledge and Discrimination. *Noûs*, 44(2), 245-268

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

RELEVANT ALTERNATIVES, PERCEPTUAL KNOWLEDGE AND DISCRIMINATION

ABSTRACT. This paper examines the relationship between perceptual knowledge and discrimination in the light of the so-called ‘relevant alternatives’ intuition. It begins by outlining an intuitive relevant alternatives account of perceptual knowledge which incorporates the insight that there is a close connection between perceptual knowledge and the possession of relevant discriminatory abilities. It is argued, however, that in order to resolve certain problems that face this view, it is essential to recognise an important distinction between favouring and discriminating epistemic support that is often overlooked in the literature. This distinction complicates the story regarding how an alternative becomes relevant, and in doing so weakens the connection between perceptual knowledge and discrimination. The theory that results, however—what I term a ‘two-tiered’ relevant alternatives theory of perceptual knowledge—accommodates many of our intuitions about perceptual knowledge and so avoids the revisionism of some recent proposals in the epistemological literature.

1. THE RELEVANT ALTERNATIVES ACCOUNT OF PERCEPTUAL KNOWLEDGE

Intuitively, there is a very close connection between perceptual knowledge and discrimination. Looking out of my window, I come to know that the creature before me is a goldfinch because I can discriminate goldfinches from other things that might plausibly be in the neighbourhood (such as woodpeckers, to use J. L. Austin’s (1961) example). Of course, in my present situation I can’t discriminate between goldfinches and, say, hologram goldfinches, but intuitively this sort of contrast is by-the-by. My knowing that there is a goldfinch in the garden is essentially constituted—at least in major part—by my being able to discriminate between goldfinches and plausible non-goldfinch alternatives; not by my being able to discriminate between goldfinches and *implausible* non-goldfinch alternatives, such as the ‘hologram goldfinch’ alternative.¹

This way of thinking about the relationship between perceptual knowledge and discrimination fits very neatly with an intuition that is widespread in epistemology, what I will refer to as the *core relevant alternatives intuition*. This intuition states that in order to know a proposition, *p*, what is required is that one is able to rule out all those not-*p* alternatives that are (in some sense to be specified) *relevant*. What is not required is that one is able to rule out the *irrelevant* alternatives.

Notice that the conception of perceptual knowledge just described, which treats such knowledge as being essentially concerned with the possession of certain discriminatory capacities, seems to represent one way of fleshing-out the core relevant alternatives intuition in the perceptual case. To begin with, it offers an answer to the question of what it means to ‘rule out’ an alternative. In the perceptual case at least, to be able to rule out an alternative is to be able to make the relevant discriminations between the target object and the object at issue in the alternative—e.g., to be able to discriminate between goldfinches and woodpeckers.

Moreover, we also have an answer to the crucial question of what determines relevance, at least in the perceptual case. For what makes an alternative relevant is whether it is the kind of alternative that might ordinarily obtain in one’s neighbourhood. The possibility that one is looking at a woodpecker right now is relevant to one’s belief that one is presently looking at a goldfinch because woodpeckers are the kinds of things that one might ordinarily find in one’s environment. In contrast, the possibility that one is looking at a hologram goldfinch just now is *not* relevant because this is not the kind of thing that one might ordinarily find in one’s environment.

I think the best way of capturing what is going on here is in modal terms by saying that the class of relevant alternatives is, roughly, all those alternatives that obtain in near-by possible worlds.² On this reading of relevance, if one had the misfortune to be in an abnormal environment in which there were near-by possible worlds in which what one was looking at just now was not a goldfinch but a hologram goldfinch, then in order to be able to know that what one is looking at is a goldfinch, one would have to be able to discriminate between goldfinches and hologram goldfinches. In this environment, then, it is very hard to know that one is looking at a goldfinch (in that it would require further inquiries on the part of the agent), even though this is something that is very easy to know in normal environments.

That one’s environment can have an impact on how hard it is to acquire perceptual knowledge is perfectly in accordance with intuition, however, as the famous barn façade case illustrates. Coming to know that what one is looking at is a barn in a normal environment is very easy, since it merely demands very mundane discriminatory powers. In contrast, coming to know that what one is looking at is a barn in an abnormal environment in which barn façades are the norm is very hard, since it demands very specialised discriminatory powers—in particular, it demands the ability to discriminate between barns and barn façades.³

We thus get the following relevant alternatives account of perceptual knowledge:

The Relevant Alternatives Account of Perceptual Knowledge

S has perceptual knowledge that *p* only if *S* can discriminate the target object at issue in *p* from the objects at issue in relevant alternative (not-*p*) propositions, where a relevant alternative is an alternative that obtains in a near-by possible world.⁴

This relevant alternatives account of perceptual knowledge leads to a pleasing result. Perceptual knowledge, it turns out, is often very easy to possess. At least in normal environments, I can know relatively mundane perceptual truths simply in virtue of being able to exercise relatively mundane perceptual discriminations.⁵ And this seems just about right.

2. RELEVANT ALTERNATIVES AND CLOSURE

As Fred Dretske (1970) famously showed, however, there is a problem lying in wait for any relevant alternatives account of perceptual knowledge of this sort; a problem that is brought out by considering the principle that knowledge is closed under competent deductions, or the ‘closure’ principle for short. This principle can be formulated as follows:

The Closure Principle

If *S* knows that *p*, and *S* competently deduces *q* from *p* (thereby coming to believe *q* while retaining her knowledge that *p*), then *S* knows that *q*.⁶

So construed, the principle seems utterly uncontentious. The problem posed by this principle for the relevant alternatives account of perceptual knowledge is that we only need to suppose that the agent concerned knows certain entailments, and makes competent deductions on the basis of this knowledge, in order to get a situation in which in order to know a proposition—e.g., that there is a goldfinch in one’s garden—one needs to know that certain intuitively far-fetched error-possibilities—e.g., that there is a hologram goldfinch in one’s garden—are false.

For example, suppose that one knows that one is looking at a goldfinch right now, and one also knows that if one is looking at a goldfinch right now then one is not looking at a hologram goldfinch, and one makes a competent deduction on this basis. It follows, given closure, that one knows that one is not now looking at a hologram goldfinch, and thus knowing that one is looking at a goldfinch in this case entails knowing that one is not looking at a hologram goldfinch. In effect, then, knowing that one is looking at a goldfinch in this case entails knowing that one is looking at a goldfinch *rather than* a

hologram goldfinch. The trouble is, of course, that one is unable to distinguish between goldfinches and hologram goldfinches—it is not as if one has made any special checks in this regard, for example—and yet it now seems that it is incumbent upon one to be able to rule out this possibility—i.e., know it to be false—in order to know something so mundane as that one is presently looking at a goldfinch. Perceptual knowledge of mundane truths can thus sometimes be very difficult to come by, even in normal environments. The key advantage of the relevant alternatives account of perceptual knowledge—that it can account for how perceptual knowledge can be quite easy to acquire in normal environments—therefore appears to be under threat.

Famously, Dretske concluded on this basis that closure should be abandoned. If we stick to the core relevant alternatives intuition, argues Dretske, then we should insist that knowing that one is looking at a goldfinch should never require the ability to rule out the irrelevant ‘hologram goldfinch’ alternative. Since accepting closure would require one in certain cases to rule out (i.e., know to be false) this hypothesis, then closure has to go, at least if scepticism is to be avoided.

The example that Dretske used to illustrate this point was the famous ‘zebra’ case. Imagine a person—we’ll call her Zula—who is at the zoo, and who gets a good look at one of the zebras in the clearly marked zebra enclosure. Zula has all the usual cognitive abilities and background knowledge one would expect of a normal person, and the circumstances are in every relevant respect entirely normal too. Does Zula know that what she is looking at is a zebra? Intuitively, we would say so. Her belief meets many of the criteria we might wish to lay down on a theory of knowledge. For example, it is reliably formed, it is virtuously formed, it is safe (in all near-by possible worlds in which she believes what she does in the same way as in the actual world, her belief continues to be true), it is sensitive (in the nearest possible world in which what she believes is false, she no longer believes it on the same basis as in the actual world), it is evidentially well-founded, and so on.

The trouble is, we can stipulate that Zula happens to know that if what she is looking at is a zebra, then it follows that what she is looking at is not a cleverly disguised mule, and makes a competent deduction on this basis. Given closure, it therefore follows that Zula knows that she is not looking at a cleverly disguised mule, and thus that she knows that what she is looking at is a zebra *rather than* a cleverly disguised mule. The problem, however, is that Zula is just a normal person with normal epistemic powers. Accordingly, she’s in no position to discriminate between zebras and cleverly disguised mules. It is not as if she has some special expertise in this regard—such as might be possessed by a zoologist, for example—or that she has made any special checks—such as going up to

the creature and checking for paint. Thus, closure seems to require that in order for Zula to know that she is looking at a zebra she must be able to rule out (i.e., know to be false) the cleverly disguised mule hypothesis. But since Zula lacks the discriminative abilities to do this, we have a problem. As with the goldfinch case described earlier, there is a tension between the core relevant alternatives intuition (as it is applied to perceptual knowledge) and the closure principle, and in the light of this tension Dretske's recommendation is to abandon the closure principle.⁷

Denying closure is not an easy thing to do, however, since the principle is so incredibly compelling. Moreover, I think there has been a general consensus that the kind of epistemology that Dretske and others have advanced (e.g., Nozick 1981) in order to undergird a new epistemology without closure brings with it some fairly serious problems. What is common to these views is a commitment to something like the sensitivity principle as a condition on knowledge: that one's belief should be such that, had what one believed not been true, one would not have believed it. One key problem that faces such view is that once the sensitivity principle is expressed in the right way, then it is no longer obvious that it generates the kinds of counterexamples to closure that it was designed to explain, and nor is it obvious that it is authentic to the core relevant alternatives intuition that motivated the rejection of closure in the first place.⁸

Although few these days are inclined to deny closure, there is a new view on the scene which does something very similar. This position is known as contrastivism, and has been defended by Jonathan Schaffer (2005), amongst others.⁹ According to the contrastivist, knowledge is to always be understood contrastively, in the sense that one never knows *p simpliciter*; instead, one knows *p* rather than each one of a set of contrasts (i.e., alternatives) to *p*, where knowing a proposition obtains rather than one of the contrasts is explicitly understood in terms of discriminating the target proposition from the specified contrasts. So Zula knows that what she is looking at is a zebra rather than an undisguised horse or baboon (because she can make the relevant discriminations), but she doesn't know that what she is looking at is a zebra rather than a cleverly disguised mule (because she cannot make the relevant discriminations). Crucially, though, there is no sense to the idea that Zula either does or does not know that what she is looking at is a zebra where this is not qualified to a contrast set.

Although contrastivists would claim that they do not deny closure—they would argue that they can retain a 'contrastivised' version of this principle, though I am somewhat sceptical about this—they do (at least) deny that there is any sense in which Zula knows that what she is looking at is not a

cleverly disguised mule, and in this regard their view is in the same revisionistic spirit as Dretske's non-closure view.¹⁰

There is also a second, and I think more important, sense in which the contrastivist view and the Dretskean view are closely connected. This concerns the fact that they both regard what it takes to rule out an alternative as ultimately being a discriminative capacity. For the contrastivist, this point is explicit to the view. To rule out an alternative to what one knows is to know that the target proposition obtains rather than the alternative, where knowing that one proposition obtains rather than an alternative is in turn understood in terms of the possession of the relevant discriminatory capacity. For the Dretskean, this point is a little more implicit to the view. To rule out an alternative to what one knows by the lights of this proposal is to know that the target proposition is true and that the (known) alternative is false. Crucially, however, it is clear that Dretske holds that, at least in the perceptual case, knowing that the target proposition obtains rather than a known alternative entails that one can discriminate between the object at issue in the target proposition and the object at issue in the alternative. It is only if Dretske holds this further thesis that it follows from the fact that Zula lacks the relevant discriminatory abilities that closure must fail. (We will come back to this point in a moment).

It is not peculiar to the Dretskean and contrastivist views that they ultimately understand what it is to rule out an alternative in terms of discriminatory abilities. After all, this sort of account of what it takes to rule out an alternative is implicit in the intuitive picture of perceptual knowledge that we noted at the outset of this paper, a picture on which there is a very tight connection between possessing perceptual knowledge and possessing the relevant discriminatory abilities. I suggest, however, that this conception of what it takes to rule out an alternative—and the picture of the relationship between perceptual knowledge and discrimination that goes hand-in-hand with it, including the relevant alternatives account of perceptual knowledge that this picture gives rise to—is not quite right. In particular, I will be arguing that there is a sense in which one can rule out an alternative where this is not to be construed in terms of a discriminative ability. As we will see, this point has a number of important ramifications for our understanding of knowledge, and perceptual knowledge in particular.

3. THREE PRINCIPLES: DISCRIMINATION, TRANSMISSION, FAVOURING

Let us look again at the problem that Dretske thought was posed by the closure principle. We have already noted that it is essential to the setting-up of this problem that we take for granted the kind of epistemological picture of the relationship between perceptual knowledge and discrimination that was sketched in §1. It is worthwhile spelling this point out in more detail.

Consider the following plausible principle:

The Discrimination Principle

If S has perceptual knowledge that p , and S knows that another (known to be inconsistent) alternative q does not obtain, then S must be able to discriminate between the object at issue in p and the object at issue in q .

So, for example, if Zula knows that what she is looking at is a zebra, and she also knows that what she is looking at is not a cleverly disguised mule (bearing in mind that she knows the relevant entailment)—i.e., if she knows that what she is looking at is a zebra *rather than* a cleverly disguised mule—then she must be able to discriminate between zebras and cleverly disguised mules. Dretske certainly seems to be buying into a principle of this sort, since it is only with this principle in play that it follows immediately from closure that there is a problem in the zebra case. That is, it is only if we construe the fact that Zula is able to know that what she is looking at is a zebra rather than a cleverly disguised mule as entailing a capacity to discriminate between zebras and cleverly disguised mules—something which, *ex hypothesi*, she is unable to do—that we get a straightforward contradiction. Without this principle, all that follows is that Zula is, it seems, able to know that one proposition has obtained rather than an alternative even while being unable to discriminate between the object at issue in the target proposition and the object at issue in the alternative. This in itself is mysterious, especially when we are dealing with perceptual knowledge, but it is not yet contradictory.¹¹

Suppose that we took it for granted that closure should be endorsed. We might thus regard the discrimination principle as suspect, despite its initial plausibility. The goal would then be to find a suitable explanation for the failure of the discrimination principle while also accounting for the mysterious nature of Zula's knowledge: how could it be that she knows that what she is looking at is a zebra rather than a cleverly disguised mule, given that she is unable to discriminate between zebras and cleverly disguised mules?

One way of responding to this problem could be to argue for the rejection of the discrimination principle on anti-luck grounds. After all, given Zula's epistemic position (in particular, that she knows that what she is looking at is a zebra, and that if it is a zebra then it is not a cleverly disguised mule), it is not a matter of *luck* that Zula's belief that she is not looking at a cleverly disguised mule is true. Her belief could not, for example, have very easily been false. So if knowledge is essentially non-lucky true belief, as some have argued, then there is *something* that can be said in favour of the idea that Zula has knowledge of this deduced proposition, even though she is unable to discriminate between zebras and cleverly disguised mules.¹²

Still, the nagging worry remains that it is odd that Zula is able to know such a thing given that she lacks the relevant discriminatory capacities. Indeed, even if we set aside this problem which is posed by the discrimination principle, there appears to be a further difficulty which faces the idea that Zula can know the deduced proposition. This aspect of the problem concerns the *evidential* status of Zula's putative knowledge in this respect. In particular, closure forces us to regard Zula as knowing that what she is looking at is not a cleverly disguised mule, and yet she doesn't appear to have any good supporting evidence for this knowledge. Moreover, although we might not claim that all knowledge must be evidentially grounded, all will surely agree that knowledge of a proposition like this must be so grounded.

There are two ways of bringing this evidential point into sharp relief. Perhaps the most immediate way to do this is by appealing to a version of the transmission principle, which many have distinguished from the closure principle.¹³ Now the transmission principle, as it is usually formulated, tries to do three things all at once—to capture a sense in which the way in which one knows transfers across a competent deduction; to capture a sense in which one's evidence transfers across a competent deduction; and to capture a sense in which an argument has persuasive force. These three aspects of the transmission principle are all closely related of course, but for our purposes I want to focus on simply the evidential aspect of transmission. Moreover, given our interests here, our focus will be not simply be on evidential transfer, but more specifically on the transfer of knowledge-supporting evidence.

We can formulate such an evidential version of transmission as follows:

The Evidential Transmission Principle

If the evidential support for *S*'s perceptual knowledge that *p* is evidence set *E*, and *S* competently deduces *q* from *p* (thereby coming to believe *q* while retaining her knowledge that *p*), then *E* is sufficient to be the evidential support for *S*'s knowledge that *q*.¹⁴

This principle certainly does seem compelling, at least insofar as we restrict our attention to those propositions which, if known at all, are known in virtue of appropriate supporting evidence. How could it be that you have adequate supporting evidence for perceptual knowledge of one proposition, and then undertake a competent deduction, thereby gaining knowledge of a second proposition, and yet fail to have sufficient evidence to support knowledge of the second proposition?

The problem is, of course, that when applied to the zebra case this principle entails that the evidence that Zula has in support of her knowledge that what she is looking at is a zebra ought to be sufficient to support her knowledge that what she is looking at is not a cleverly disguised mule. What is odd about this is that Zula's evidence for thinking that what she is looking at is a zebra does not seem to speak at all to the cleverly disguised mule alternative. As we noted above, it is not as if Zula has any special expertise in this regard—such as might be possessed by a zoologist, for example—nor has Zula made any special checks. It is thus mysterious how her evidence could transfer across the deduction. One who wishes to retain closure thus seems committed to denying not just the discrimination principle, but also the evidential transmission principle.

The second way to bring the evidential problem into sharp relief—indeed, a way of bringing the problem into sharp relief that gets more to the heart of the matter I think—is by appeal to what is known as the favouring principle.¹⁵ We can formulate this principle as follows, where 'better' evidence is evidence that makes the target proposition more likely to be true:

The Favouring Principle

If S has perceptual knowledge that p and S knows that q (and knows that p entails q), then S has better evidence in support of her belief in p than for believing not- q .

As with the evidential transmission principle, this principle certainly seems compelling, at least as regards those propositions which, if known at all, one knows in virtue of possessing appropriate supporting evidence. How can it be that one knows that one proposition obtains and that a second, known to incompatible, alternative does not obtain, and yet one lacks better evidence for believing the first proposition rather than believing the alternative? Put more simply, how can one know that one alternative obtains rather than a second alternative, when one lacks better evidence for thinking that the first alternative obtains rather than the second?

Applied to the zebra case, however, this principle entails that if Zula knows that what she is looking at is a zebra, and also knows that what she is looking at is not a cleverly disguised mule (and the relevant entailment), then she must have better evidence for believing that what she is looking at is

a zebra than for believing that what she is looking at is a cleverly disguised mule. Such a demand seems entirely reasonable, but the problem, of course, is that intuitively Zula lacks such ‘favouring’ evidence in support of her belief. Intuitively, that is, since her evidence does not speak to the cleverly disguised mule alternative it can hardly be thought to favour her believing that what she is looking at is a zebra over her believing the alternative that what she is looking at is a cleverly disguised mule. Anyone retaining closure in the zebra case therefore seems to be required to deny this principle as well. The problems facing any view which retains closure thus seem to be mounting up.¹⁶

So if we wish to retain the closure principle then there is a burden upon us to explain what this knowledge of the entailed proposition is in virtue of, and that means that we both need to account for the failure of the discrimination principle whilst also responding to this evidential challenge posed by the evidential transmission and favouring principles. I think that the key to resolving both these problems is to recognise that Zula’s evidential position is in fact much stronger than we ordinarily tend to think, and certainly strong enough to satisfy the requirements laid down by the evidential transmission and favouring principles.

4. FAVOURING AND DISCRIMINATING EPISTEMIC SUPPORT

Consider again the evidential support that Zula has for her knowledge that what she is looking at is a zebra. Now one might naturally say that Zula’s knowledge of this proposition is knowledge that she has gained just by looking, and this way of speaking implies that the only evidence that Zula has in favour of her belief is the evidence she gets just by looking—i.e., evidence regarding the visual scene before her. If this is all the evidence that Zula has for her belief—and granting that she has knowledge of the relevant entailment—then it ought to be clear that her knowledge of this proposition fails to satisfy the favouring principle, since the evidence that she gains just by looking in no way favours the ‘zebra’ alternative over the ‘cleverly disguised mule’ alternative. And if it fails to satisfy the favouring principle, then how can it satisfy the evidential transmission principle?

I’m actually a little suspicious of the idea that Zula can come to know that what she is looking at is a zebra ‘just by looking’, at least if this phrase is meant to indicate that the evidential support that Zula has for her belief is merely that evidence which is offered by the bare visual scene before her. That is, I think that while there is a sense in which it is obviously true that Zula gains her knowledge

just by looking—in that Zula does nothing more than look in order to gain her knowledge—one can grant this straightforward reading of this locution while objecting to the evidential conclusion that is drawn from it. After all, perceptual knowledge can at least sometimes—and perhaps often, or even always—involve a wide range of specialist expertise and background knowledge. Without such expertise and background knowledge, one may look all one wants and still come to know nothing of consequence. But given that such expertise and background knowledge would surely have ramifications for the total evidence that you possess in support of your belief, it follows that coming to know a proposition just by looking need not entail that the only evidence you possess for your belief is the evidence you gained from the bare visual scene before you.

We will set this concern to one side, however, since even if one grants that someone can come to know a proposition on the basis of evidence gained merely by looking in this restrictive sense, it still remains that there is a problem with the standard way of understanding the evidential basis of Zula's belief that what she is looking at is a zebra. Let us grant, then, that Zula's belief in this proposition is adequately supported by the evidence she gains from the bare visual scene presented to her, no matter what further evidence she might possess in favour of this belief. Keeping this set of evidence fixed, and now imagining Zula becoming aware of the relevant entailment and making a competent deduction on this basis, we are faced with the question of how Zula's evidence, so construed, could in any way be good evidence for believing the deduced proposition. How can the evidence Zula gains just by looking in this case have any bearing on the possibility that what she is looking at is a cleverly disguised mule? More pertinently, how can the evidence that she gains in this way supply her with better reason to believe that what she is looking at is a zebra rather than a cleverly disguised mule, as the favouring principle would demand? Moreover, how can the evidence which she gains in this way supply her with knowledge-supporting evidence in favour of her deduced belief that she is not looking at a cleverly disguised mule, as the evidential transmission principle would demand?

Notice, however, that in the setting-up of this problem a key issue is being glossed over, which is what epistemic effect Zula's becoming aware of the entailment has. It is surely right to suppose that Zula can come to know that what she is looking at is a zebra without having any awareness of the cleverly disguised mule hypothesis. But is it really so plausible that Zula can become aware of this error-possibility and its incompatibility with what she believes, *and* retain her knowledge that what she is looking at is a zebra, even while taking no view at all regarding what would entitle her to dismiss such an error-possibility?

In order to see the implausibility of this suggestion, put yourself into Zula's shoes for a moment. Surely in becoming aware of this error-possibility and its incompatibility with what you believe, you would form a view about what entitles you to dismiss this possibility (assuming, of course, that you retain your original belief, and so think that you should dismiss this possibility)? Moreover, I take it that there is quite a lot of evidence that you can offer in favour of taking this view. Think, for example, of the wealth of background knowledge that you have which is relevant to this error-possibility—in particular, evidence regarding the likelihood of this error-possibility obtaining. One might reason, for instance, that there would be no point in such a deception, that it would be costly and time-consuming without bringing any comparable benefit, that it would be easily found out, and then the zoo-owner would be subject to penalties, and so on. What is important about this process of taking a view about one's entitlement to dismiss this error-possibility is that in engaging in this process one thereby highlights that one has better evidence for believing that one is presently looking at a zebra than for the alternative that one is looking at a cleverly disguised mule.

The favouring principle is thus met in this case, because one does have better evidence for believing that what one is looking at is a zebra rather than a cleverly disguised mule. Moreover, since such favouring evidence is possessed there is surely no problem with supposing that the evidential transmission principle is met as well. The point is that in order to undertake the competent deduction at issue in the evidential transmission principle—which, recall, is a competent deduction which preserves the agent's knowledge of the antecedent proposition—it is essential that the agent's evidence set incorporates the favouring evidence that the agent has for preferring the 'zebra' alternative over the 'cleverly disguised mule' alternative. But if the agent has supporting evidence of *this* sort, then it is no longer mysterious that she has evidentially supported knowledge of the deduced proposition.

Of course, it is always possible that Zula lacks the kind of background knowledge at issue here, and so has no good reason to dismiss this error-possibility. But notice that if this is the case then it is no longer plausible to suppose that she retains her knowledge of the target proposition. After all, it is one thing to grant that Zula can have knowledge of this proposition while having no awareness of this error-possibility, and quite another to suppose that she retains this knowledge even while being aware of this error-possibility and being unable to rationally dismiss it. That is, becoming aware of an error-possibility that you know is incompatible with what you believe and being unable to rationally dismiss it is, I would claim, knowledge-defeating.¹⁷

This is a good point to review where we are. To begin with, we have argued that insofar as Zula retains her perceptual knowledge that what she is looking at is a zebra then it must be the case that she is in a position to rationally dismiss the cleverly disguised mule error-possibility once she becomes aware of it. Crucially, however, we have also seen that it is perfectly consistent with the idea that Zula is unable to discriminate between zebras and cleverly disguised mules that she is nevertheless in possession of evidence which would enable her to rationally dismiss this error-possibility, thereby satisfying both the favouring principle and the evidential transmission principle. In such a case, her belief that what she is looking at is not a cleverly disguised mule is appropriately evidentially grounded after all, even though Zula cannot make the relevant discrimination. The upshot of this is that one can indeed have *bona fide* evidentially supported knowledge that one alternative has obtained rather than another (known to be incompatible) alternative, without this requiring that one has the relevant discriminative capacity. The discrimination principle is thus too strong, and must be rejected. Moreover, notice that we have achieved this result while appealing to our ordinary intuitive conception of evidence. This result is therefore not hostage to a contentious account of evidence.¹⁸

A distinction has thus opened up between the kind of epistemic support provided by favouring evidence, and the kind of epistemic support provided by discriminatory capacities, since we have seen that one can have knowledge-supporting favouring evidence even while lacking the relevant discriminatory capacities. In failing to recognise this distinction, the participants in this debate have unduly limited their dialectical options. It is perfectly compatible with the idea that Zula cannot discriminate between zebras and cleverly disguised mules that she nevertheless has the relevant knowledge-supporting evidence for her beliefs so as to satisfy the favouring and transmission principles. One can therefore accept, in line with closure, that Zula can come to know that she is not looking at a cleverly disguised mule without thereby becoming puzzled, given Zula's limited discriminatory capacities, as to how such knowledge is possible.¹⁹

5. DIAGNOSIS

A question that one might naturally ask at this point is why this distinction between favouring and discriminating epistemic support has not been widely recognised, especially given that I claim that it is rooted in our everyday conception of the epistemological landscape (and so is not hostage to a specific

view of evidence). I think that there are several explanations for this. One explanation concerns the point noted above that when we think about the zebra case we tend to naturally describe Zula's knowledge as knowledge that is gained 'just by looking'. In doing so, we have a tendency to ignore the background knowledge that can be relevant to the evidential support that an agent has for her perceptual knowledge, even in cases where the belief formed is fairly mundane.

Interestingly, even when commentators do consider the wider evidential standing of Zula's belief they still understand such evidence in an unduly restrictive way. For example, in his discussion of Zula's epistemic position, Dretske (1970, 1016) asks whether Zula has examined "the animals closely enough to detect such a fraud", and it is clear from the ensuing remarks that he thinks that a negative answer to this question has a fairly decisive implication for whether or not we should ascribe knowledge to Zula. Such grounds would indeed be very useful for Zula to have in this regard, but the lack of them does not indicate a lack of knowledge since, as we have seen, it remains that Zula may have adequate favouring evidence in support of her beliefs. Notice that what is happening here is that the kind of additional evidence which Dretske is admitting as relevant is only that *discriminating evidence* which suggests that Zula is able to make the relevant discrimination (since if she had made such special checks, then she would have been in a position to discriminate between a zebra and a cleverly disguised mule in this case). One can have the relevant favouring evidence, however, even while lacking such discriminating evidence.

Many commentators have followed Dretske in treating Zula's evidence in this way. Consider the following passage from a recent paper by Crispin Wright:

You go to the zoo, see several zebras in a pen, and opine (ZEBRA): that those animals are zebras. Well, you know what zebras look like, and these animals look just like that. Surely you are fully warranted in your belief. But if the animals are zebras, then it follows that they are not mules painstakingly and skilfully disguised as zebras. Does your warrant transmit to the latter claim? Did you examine the animals closely enough to detect such a fraud? Almost certainly not. The grounds you have for (ZEBRA)—essentially, just the look of the beasts—have no bearing on this possibility. (Wright 2003, 60)

Like Dretske, Wright also assumes that the only additional evidence that would be relevant would be discriminating evidence, thereby ignoring the possibility of favouring evidence. In doing so he offers an account of this case in which the agent concerned has knowledge of the target proposition and yet has an epistemic position which fails to satisfy either the favouring or evidential transmission principles.²⁰ Moreover, notice that Wright in addition tends to identify the evidential support that the agent has for her belief with merely the evidence she gains from the bare visual scene presented to

her—her evidence is, “essentially, just the look of the beasts”—when, as we have seen, this is also an unduly restrictive way of understanding the agent’s evidential standing. While we have granted for the sake of argument that Zula might well have knowledge that what she is looking at is a zebra merely on this bare evidential basis, it does not follow that Zula can retain this knowledge merely on this basis once she becomes aware of certain error-possibilities.

A further reason why commentators have tended to overlook this distinction between favouring epistemic support and discriminating epistemic support is that, at least typically, when one claims to know a proposition one thereby represents oneself as possessing discriminating evidence in support of that claim rather than just favouring evidence. For example, if Zula were to flatly (i.e., without qualification) claim to know that the creature before her is a zebra in a conversational context in which the cleverly disguised mule error-possibility is at issue, then we would surely regard her as implying that she has discriminating evidence in favour of her belief, such as that she has some special expertise or knowledge in this regard (that she’s made special checks, say). If she is unable to do this, then we would expect her to indicate this fact. She might qualify her claim to know, for example, by specifying that the only grounds she has in support of her assertion are favouring grounds.

It is an interesting question just why appropriate (unqualified) claims to know typically generate this sort of implicature, though this is not an issue that we can usefully engage with here. What is important about this observation, however, is that it highlights one reason why one might naturally focus on whether Zula has discriminatory epistemic support when considering her epistemic standing, to the exclusion of other types of epistemic support, especially favouring epistemic support.²¹

One final reason why I think the distinction between favouring and discriminating epistemic support is often overlooked when it comes to these cases is that discussion of zebra-style examples often takes place with one eye on the radical sceptical problem. On the face of it, this might seem entirely understandable, since the sceptical challenge—as it is usually formulated at any rate—does seem to be a type of zebra-style case. For example, suppose that one has perceptual knowledge that one has hands. Suppose further that one knows that if one has hands then one is not a (handless) brain-in-a-vat (BIV), and one undertakes a competent deduction on this basis. Given the closure principle, then, one knows that one is not a BIV. But how can that be, given that, *ex hypothesi*, one cannot discriminate between having hands and being a BIV? Moreover, the evidential transmission principle and the favouring principle both seem to fail in this case—intuitively, one does not have better evidence for believing that one has hands than for the BIV hypothesis, and neither does the putatively

knowledge-supporting evidence one has for believing that one has hands seem to transfer across a competent deduction to be knowledge-supporting evidence for believing that one is not a BIV.

It is probably best, however, to keep the issue posed by the zebra case and that posed by the sceptical case separate. Indeed, the distinction drawn here between favouring and discriminating epistemic support indicates why, for on closer analysis the sceptical case can be shown to pose a far trickier challenge. In particular, a crucial disanalogy between the two cases is that sceptical error-possibilities by their nature call one's evidence into question *en masse*. Intuitively, if one is a BIV, then one has very little evidence for thinking that one has hands. In contrast, the truth of the cleverly disguised mule hypothesis has, on the face of it anyway, very little bearing on the strength of Zula's evidence that what she is looking at is a zebra. I think this distinctive feature of sceptical cases ensures that one cannot simply make use of the distinction between favouring and discriminating epistemic support in order to respond to the sceptical case, since it is moot whether one should allow the usual scope of evidence that one would typically ascribe to an agent in the light of a sceptical challenge.

For example, just as we appealed to Zula's background knowledge about the plausibility of the cleverly disguised mule hypothesis in order to accord her with the appropriate favouring evidence in support of her beliefs, suppose we appealed to similar background knowledge that our agent in the sceptical case might be thought to have in order to accord her with the requisite favouring evidence. After all, we might naturally suppose that the agent in the sceptical case possesses all sorts of good reasons for thinking that the BIV hypothesis is implausible. The problem, however, should now be manifest. We can unproblematically appeal to such background knowledge in the case of Zula precisely because the error-possibility at issue does not call into question this background knowledge. The same is not true of the sceptical case. If the agent in the sceptical case is indeed a BIV, then she would lack such background knowledge. Accordingly, it would be contentious to appeal to such evidence as a means of showing that this agent's beliefs have the required epistemic support. Given this crucial disanalogy between zebra-style cases and the sceptical case, having the sceptical problem in the background can thus tend to make the evidential problem posed by zebra-style cases look more intractable than it in fact is.²²

6. A TWO-TIERED RELEVANT ALTERNATIVES THEORY

In any case, the state-of-play is that we have identified an important distinction between favouring and discriminating epistemic support, a distinction that can enable us to explain why the discrimination principle should fail while also accounting for the evidential basis of Zula's knowledge that what she is looking at is not a cleverly disguised mule. One might ask where this leaves the relevant alternatives account of perceptual knowledge that we began with.

Recall that this account of perceptual knowledge held that perceptual knowledge is essentially constituted in terms of the possession of those capacities necessary to discriminate between the object at issue in the target proposition and the objects at issue in the relevant alternatives, where the relevance of an alternative was in turn determined by one's environment. It should be clear from the foregoing that we cannot maintain this conception of perceptual knowledge in its current form since we have accepted that it is a consequence of closure that in order to have perceptual knowledge of a proposition it is sometimes necessary to possess not just the discriminatory capacities just described, but also relevant favouring evidence. Although this conclusion means that we need to reject the relevant alternatives account of perceptual knowledge as it presently stands, I think we can nevertheless retain the spirit of the view. In particular, I suggest that what we need to do is replace the relevant alternatives account of perceptual knowledge with a two-tiered view that is cast along the same lines. Let me explain.

To begin with, we need to distinguish between a narrow and a broad conception of relevance. The narrow conception of relevance is simply that notion at issue in the relevant alternatives account of perceptual knowledge which concerns those error-possibilities that obtain in near-by possible worlds. One thing that is absolutely right about the relevant alternatives theory of perceptual knowledge is that when it comes to alternatives which are narrowly relevant the possession of the relevant discriminatory capacity *is* essential for knowledge. So if Zula is indeed in an environment in which she could very easily be looking at a cleverly disguised mule just now, then in order to know that the creature before her is a zebra she must be able to discriminate between zebras and cleverly disguised mules. Since she cannot do this, in this environment she lacks knowledge that what she is looking at is a zebra. But that is, I would argue, entirely what we would expect.

Thus far, then, I am agreeing with Dretske and many others who follow Dretske in holding that there is a very tight connection between perceptual knowledge and discrimination. We diverge when it

comes to those non-narrow alternatives which are also, I would claim, relevant. These are alternatives which are not narrowly relevant in the sense just specified, but which are made relevant by other considerations. One way in which this can happen—and there may be other ways, though I won't be taking a stand on this issue here—is when one becomes aware of an error-possibility which is incompatible with something which one believes. I have argued that if one is to retain one's knowledge of the original proposition, then one must be able to rule out this error-possibility.

On the face of it, admitting this much would seem to entail that in order to have perceptual knowledge of quite mundane propositions, such as that the creature before one is a zebra, one has to know the denials of some quite extraordinary error-possibilities, such as that the creature before one is not a cleverly disguised mule. Accordingly, one might think that perceptual knowledge of quite mundane propositions becomes unduly difficult, and therefore doubt the epistemology that lies behind this demand. As we have seen, however, the epistemic demands made by this requirement are actually quite modest, since all one needs to satisfy it is favouring supporting evidence, and this is usually quite easy to come by, at least if one is reasonably intellectually sophisticated. What awareness of these non-narrowly relevant alternatives does not demand, crucially, is the sort of discriminatory abilities that *would* be rather hard to come by, and it is only if that were the case that this thesis would make perceptual knowledge of the relatively mundane unduly difficult to acquire. And since one can usually straightforwardly satisfy this requirement, then accepting such epistemic principles as closure, favouring and evidential transmission no longer poses a problem because in accepting them one is not obliged to thereby accept that perceptual knowledge of the relatively mundane can sometimes entail quite extraordinary discriminatory powers that one is unlikely to possess.

We thus get a two-tiered picture of relevance which is very much in accord with commonsense. For one thing, it confirms to our intuitive picture of the relationship between perceptual knowledge and discrimination that we began with, since there *is* still a very close relationship between these two notions in play here. Moreover, it conforms to the basic thought that underlies the core relevant alternatives intuition, which is that one can have knowledge without having to rule out far-fetched error-possibilities. This is still true. If one is not aware of the non-narrow alternatives—and such alternatives are not made non-narrowly relevant in some other way—then one can know that what one is looking at is a goldfinch without having to rule out the possibility that what one is looking at is a hologram goldfinch. Crucially, however, this two-tiered relevant alternatives theory also accommodates the thought that alternatives can be made relevant in other ways than by one's

environment, though it does so in such a way as to keep the demands on our discriminative powers realistic.²³

That becoming aware of an alternative can make that alternative relevant explains, I think, why raising an alternative in a conversational context can make evidential demands on one. Suppose I claim to know that the creature over there is a goldfinch, and you respond by introducing the suggestion that it might be a hologram goldfinch (to keep the example as clean as possible, let us stipulate that you don't offer any grounds for thinking that this alternative might obtain; you just simply mention it as an alternative). I may well tell you to get lost, and we may well grant that I would be right to do so (you've certainly offended against some normal constraints on good conversation). Still, if I hadn't considered the possibility before, I should now, since I now know that this is an alternative to what I take myself to know. Moreover, in considering the alternative, I should be able to rule it out. If ruling it out meant having the required discriminative capacity, then knowledge would be in short supply with people like you around, but if it merely means having appropriate favouring evidence, then meeting this requirement should be fairly straightforward. I have all kinds of good reasons for thinking that what I am looking at is a goldfinch rather than a hologram goldfinch.²⁴

This simple distinction regarding how alternatives can become relevant—and the distinction between favouring and discriminating epistemic support that underlies it—can, I think, remove much of the impetus for one key thread of revisionistic thought in epistemology during the last thirty or so years since Dretske published his seminal article. One aspect of this thread of revisionistic thought is the Dretskean thesis that the principle of closure must be rejected, but it is not the only aspect. Now is not the place to explore this claim fully, but I will close by offering a 'taster' of what I have in mind.

Consider again the contrastivist account of knowledge described earlier. As we saw above, according to this proposal knowledge is to be understood as essentially involving discrimination, such that knowing a proposition boils down to having the relevant discriminatory capacities. The distinction between favouring and discriminating epistemic support clearly removes a central plank of support for this view. After all, one principal motivation for the position comes from the intuition that it is problematic to suppose that Zula is able to know that what she is looking at is not a cleverly disguised mule (and thereby know that what she is looking at is a zebra rather than a cleverly disguised mule). But as we have demonstrated here, this supposition is only problematic provided one already understands such knowledge in terms of Zula's possession of the relevant discriminatory powers. It is

thus highly contentious to move from this ‘intuition’ to the claim that knowledge should be understood contrastively, and thereby essentially in terms of discriminatory powers.

It should be clear even from this brief discussion of contrastivism that the proposal sketched here has ramifications for any revisionist account of knowledge which takes as part of the philosophical ‘data’ that needs to be explained the fact that agents like Zula in zebra-style cases lack knowledge of the entailed proposition, for it is precisely this ‘intuition’ which is shown to be unmotivated by this proposal.²⁵ Recognising the distinction between favouring and discriminating epistemic support thus has important implications not just for our understanding of the relevant alternatives intuition and the relationship between perceptual knowledge and discrimination, but also for contemporary epistemology more generally.²⁶

REFERENCES

- Austin, J. L. (1961). ‘Other Minds’, in his *Philosophical Papers*, (eds.) J. O. Urmson & G. J. Warnock, Oxford: Clarendon Press.
- Black, T. (2002). ‘A Moorean Response to Brain-In-A-Vat Scepticism’, *Australasian Journal of Philosophy* 80, 148-63.
- Cohen, S. (1988). ‘How to be a Fallibilist’, *Philosophical Perspectives* 2, 91-123.
- (2000). ‘Contextualism and Skepticism’, *Philosophical Issues* 10, 94-107.
- Davies, M. (2004). ‘Epistemic Entitlement, Warrant Transmission and Easy Knowledge’, *Proceedings of the Aristotelian Society* 78 (supp. vol.), 213-45.
- DeRose, K. (1995). ‘Solving the Skeptical Problem’, *Philosophical Review* 104, 1-52.
- Dretske, F. (1970). ‘Epistemic Operators’, *Journal of Philosophy* 67, 1007-23.
- (2005a). ‘The Case Against Closure’, *Contemporary Debates in Epistemology*, (eds.) E. Sosa & M. Steup, 13-26, Oxford: Blackwell.
- (2005b). ‘Reply to Hawthorne’, *Contemporary Debates in Epistemology*, (eds.) E. Sosa & M. Steup, 43-6, Oxford: Blackwell.
- Goldman, A. (1976). ‘Discrimination and Perceptual Knowledge’, *Journal of Philosophy* 73, 771-91.
- Hawthorne, J. (2004). *Knowledge and Lotteries*, Oxford: Clarendon Press.
- (2005). ‘The Case for Closure’, *Contemporary Debates in Epistemology*, (eds.) E. Sosa & M. Steup, 26-43, Oxford: Blackwell.
- Klein, P. (1995). ‘Skepticism and Closure: Why the Evil Genius Argument Fails’, *Philosophical Topics* 23, 213-36.
- Lewis, D. (1996). ‘Elusive Knowledge’, *Australasian Journal of Philosophy* 74, 549-67.
- McDowell, J. (1995). ‘Knowledge and the Internal’, *Philosophy and Phenomenological Research* 55, 877-93.
- Neta, R. (2002). ‘S Knows that P’, *Noûs* 36, 663-81.

- (2003). ‘Contextualism and the Problem of the External World’, *Philosophy and Phenomenological Research* 63, 1-31.
- Nozick, R. (1981). *Philosophical Explanations*, Oxford: Oxford University Press.
- Pritchard, D. H. (2002). ‘Radical Scepticism, Epistemological Externalism, and Closure’, *Theoria* 68, 129-61.
- (2005a). ‘Contrastivism, Evidence, and Scepticism’, *Social Epistemology* (special issue on *Contrastivism*) 20.
- (2005b). *Epistemic Luck*, Oxford: Oxford University Press.
- (2005c). ‘The Structure of Sceptical Arguments’, *The Philosophical Quarterly* 55, 37-52.
- (2007a). ‘Anti-Luck Epistemology’, *Synthese* 158, 277-97.
- (2007b). ‘How to be a Neo-Moorean’, *Internalism and Externalism in Semantics and Epistemology*, (ed.) S. Goldberg, Oxford: Oxford University Press.
- (2007c). ‘McDowellian Neo-Mooreanism’, *Disjunctivism: Perception, Action, Knowledge*, (eds.) A. Haddock & F. Macpherson, Oxford: Oxford University Press.
- Schaffer, J. (2005). ‘Contrastive Knowledge’, *Oxford Studies in Epistemology*, (eds.) T. Gendler & J. Hawthorne, Oxford: Oxford University Press.
- (*Forthcoming*). ‘Closure, Contrast, and Answer’, *Philosophical Studies*.
- Sinnott-Armstrong, W. (2006). *Moral Skepticisms*, New York: Oxford University Press.
- Sosa, E. (1999). ‘How to Defeat Opposition to Moore’, *Philosophical Perspectives* 13, 141-54.
- (2000). ‘Skepticism and Contextualism’, *Philosophical Issues* 10, 1-18.
- Stanley, J. (2005). *Knowledge and Practical Interests*, Oxford: Clarendon Press.
- Stine, G. (1976). ‘Skepticism, Relevant Alternatives, and Deductive Closure’, *Philosophical Studies* 29, 249-61.
- Unger, P. (1968). ‘An Analysis of Factual Knowledge’, *Journal of Philosophy* 65, 157-70.
- Williams, M. (1991). *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*, Oxford: Blackwell.
- Williamson, T. (1997). ‘Knowledge as Evidence’, *Mind* 424, 717-41.
- (2000). *Knowledge and Its Limits*, Oxford: Oxford University Press.
- Wright, C. J. G. (2003). ‘Some Reflections on the Acquisition of Warrant by Inference’, *New Essays on Semantic Externalism and Self-Knowledge*, (ed.) S. Nuccetelli, Cambridge, MA: MIT Press.
- (2004). ‘Warrant for Nothing (and Foundations for Free)?’, *Proceedings of the Aristotelian Society* 78 (supp. vol.), 167-212.

NOTES

¹ Of course, not all perceptual knowledge is knowledge of objects, since, for example, some perceptual knowledge is of distances. I take it, however, that perceptual knowledge is paradigmatically about objects in this way, and so in order to simplify matters in what follows I will set these other types of perceptual knowledge to one side. (If one prefers, then one can think of the type of perceptual knowledge at issue in this paper as specifically *objectual* perceptual knowledge).

² Notice that I am understanding the ordering of possible worlds in the standard way in terms of their similarity to the actual world. A different way of understanding relevance is in probabilistic terms, such that what makes an alternative irrelevant is the fact that it concerns low-probability possibilities (e.g., hologram goldfinches). Although the difference between the two views is not particularly important here (because, for the most part, low probability possibilities are far-off possibilities, and *vice versa*), I favour the first sort of proposal because of the fact that low probability events can occur in near-by possible worlds (think, for example, of lottery wins). In such cases, I maintain that the modal nearness of the possibility will make it relevant even despite the fact that it is a low probability event. For a defence of a version of the probabilistic account, see Cohen (1988).

³ This example is described in Goldman (1976), and credited to Carl Ginet.

⁴ Note that one would need to supplement this principle in various ways if one wished to provide a complete account of perceptual knowledge (even granting the *caveat* noted in endnote 1 above). For example, the class of relevant alternatives will need to be extended beyond the not-*p* alternatives, since there will clearly be some alternatives which are incompatible with knowing that *p* even while being compatible with *p*. Even so, I think it is reasonable to refer to such a proposal as an ‘account’ of perceptual knowledge, even though it is incomplete in these ways, since the claim is that perceptual knowledge *essentially* consists in the satisfaction of this condition (i.e., setting some peripheral considerations to one side, if one satisfies this condition then one possesses perceptual knowledge). I am grateful to an anonymous referee for *Noûs* for pressing me on this issue.

⁵ Notice as well that this account of perceptual knowledge is very anti-intellectualist, which might also be thought to be an advantage of the proposal. For example, a small child may lack the concept of a horse and yet, because she can nevertheless discriminate between zebras and horses, she can still come to know that what she is looking at is a zebra.

⁶ This is a different—and, I think, more plausible—way of understanding the closure principle to how Dretske (1970; 2005a) understands it, though nothing hangs on this difference here. This formulation of the closure principle is essentially that offered by Williamson (2000, 117) and Hawthorne (2005, 29). For the most recent critical discussion of the closure principle, see the exchange between Dretske (2005a; 2005b) and Hawthorne (2005).

⁷ I take it as obvious that this problem isn’t peculiar to the zebra case (or the goldfinch case for that matter), since the problem straightforwardly generalises to lots of other cases. Consider Barney, who is presently looking (in normal circumstances) at a barn and, on this basis, comes to know that what he is looking at is a barn. Suppose, however, that Barney knows that if what he is looking at is a barn, then it isn’t a barn façade, and competently deduces on this basis that he is not looking at a barn façade. He thereby knows, given closure, that what he is looking at is a barn rather than a barn façade. We can stipulate, however, that Barney has no special evidence or expertise in this regard. He has not checked to see whether this is a barn façade, nor is he skilled enough to be able to tell, just by looking straight on, whether the structure before him is a barn rather than a barn façade. In short, he cannot discriminate between barns and barn façades. So how then can it be that he knows that what he is looking at is a barn rather than a barn façade, which is what closure seems to demand?

⁸ This is not the place to explore these points further. For a defence of the first claim—that the sensitivity principle, properly understood, does not generate the kinds of closure-failure that it is meant to—see Williams (1991, ch. 9) and Black (2002). For a defence of the second claim—that a sensitivity-based epistemology is in any case in conflict with the core relevant alternatives intuition—see Pritchard (2002; 2005b, chs. 2-3). There are lots of other objections that have been levelled against sensitivity-based views, of course. For a recent statement of one such objection, see Sosa (1999).

⁹ For another prominent defence of contrastivism, see Sinnott-Armstrong (2006).

¹⁰ For a contrastivist defence of closure-type inferences, see Schaffer (*forthcoming*). I discuss the contrastivist account in more detail in Pritchard (2005a).

¹¹ As Dretske exegesis, this is not quite right, since Dretske (1970, 1016) not only appeals to Zula’s lack of the relevant discriminative abilities in his argument against closure, but also maintains that she lacks the required supporting evidence for knowledge of the entailed proposition. Crucially, though, the kind of supporting evidence that Dretske thinks would be required is evidence for thinking that she can make the relevant discriminations, such as the evidence she would gain by making special checks, and so this second aspect to his critique of closure is ultimately just a variation on the first. I consider the import that such ‘discriminating’ evidence has for this issue below.

¹² The idea that knowledge is essentially non-lucky true belief is one of the recurring *motifs* of epistemology. The main stimulus for the contemporary discussion of an anti-luck epistemology of this sort is Unger (1968). For a recent defence of anti-luck epistemology—one which ties such an epistemology to the kind of safety-based accounts of knowledge advocated by Sosa (1999; 2000) and others—see Pritchard (2005b; 2007a).

¹³ For the most recent exchange on the transmission principle, see Davies (2004) and Wright (2004).

¹⁴ One complication that I will be setting aside in what follows is that one might plausibly maintain that simply undertaking the relevant deduction enhances one’s evidence set. Klein (1995), for example, holds a view of roughly this sort. If that’s right, however, then it could be that it is only one’s *enhanced* evidence set which is sufficient for knowledge of the deduced proposition in some cases. Since the evidential worry highlighted by evidential transmission that we are discussing here concerns the fact that the agent’s evidence for believing the deduced proposition is *nowhere near* sufficient for knowledge of that proposition, I think we can safely ignore this potential complication.

¹⁵ A version of this principle is also known in the literature as the ‘underdetermination’ principle. For a more detailed discussion of this principle, including its logical relationship with the closure principle, see Pritchard (2005b, ch. 4; 2005c).

¹⁶ Stine (1976) defends the closure principle by in effect denying the two evidential principles just mentioned. This is because she accommodates the oddity of allowing that an agent can know the denials of irrelevant alternatives like the cleverly disguised mule alternative by claiming that such knowledge is non-evidential. That is, while she agrees with Dretske that Zula has no evidence for believing the cleverly disguised mule hypothesis to be false, she nevertheless argues that because this is an irrelevant alternative this does not prevent her from knowing this to be the case, since one can know that irrelevant alternatives are false without evidence. The way she retains closure thus compels her to abandon the evidential and favouring principles. As I argue in the rest of the paper, however, such a concession can be avoided if one understands the evidential position of the agent in the right way.

¹⁷ What should we say about agents who ought to be aware of a certain error-possibility which they are unable to rationally dismiss but aren't (e.g., cases in which the agent goes to great lengths to ensure that she is not perturbed by such intellectual 'irritations')? I comment on this complication below in endnote 23.

¹⁸ For example, if one thought, with Williamson (1997; 2000, ch. 9), that all evidence is knowledge, then that Zula knows that what she is looking at is a zebra would itself constitute favouring evidence. Given that knowledge is factive, where Zula has this knowledge she clearly has better evidence for believing that what she is looking at is a zebra rather than a cleverly disguised mule. Indeed, any conception of evidence which regarded Zula's evidence as suitably factive would do the trick. The conception of evidence offered by McDowell (1995), for example, would treat Zula's seeing that there is a zebra before her as part of her supporting evidence, and yet one can only see that something is the case if it is the case, and so her evidence would entail what it is evidence for. Hence, it couldn't fail to thereby be favouring evidence as well. The argument I offered above doesn't make any appeal to factive evidence, however, and I think that it is all the stronger for it (though this is not of course to suggest that such a conception of factive evidence should be rejected). For an account of factive evidence which is lodged within a contextualist framework, see Neta (2002; 2003). In Pritchard (2007c) I argue that this distinction between favouring and discriminating epistemic support is vital to proponents of factive conceptions of epistemic support, since without it they will be unable to explain why an agent can have evidentially supported knowledge that one alternative has obtained rather than another alternative (e.g., a sceptical alternative) even while lacking the relevant discriminatory capacities.

¹⁹ It is important to note that this proposal is very different to a superficially similar view defended by Cohen (1988). While Cohen also makes the point that an agent like Zula might be in possession of background evidence that would enable her to rationally dismiss the cleverly disguised mule hypothesis, he makes the crucial error of supposing that being in possession of this evidence is required for anyone to come to know that what they are looking at is a zebra. But this is a far too restrictive account of knowledge, and certainly does not conform to the original aspirations of a relevant alternatives account of perceptual knowledge, a view that is meant to set the bar for perceptual knowledge potentially quite low. Imagine, for example, a small child looking at a zebra in a zoo in normal conditions (there's no deception going on, no reason to think that there is deception going on, no far-fetched error-possibilities have been raised, and so on). This child can perfectly well discriminate zebras from, say, other things that might plausibly be around (horses, litter bins, etc.). We would not expect such a child to have the kind of collateral evidence available to her that Cohen wants to make a pre-requisite of perceptual knowledge. Surely, though, the child knows that what she sees is a zebra? In any case, one of the advantages of my approach is that it is perfectly consistent with allowing knowledge in cases like this. Moreover, there is no need on my view to postulate the kind of attributer contextualism about 'knows' that Cohen endorses in the light of the zebra case. For him, what happens when we consider the cleverly disguised mule hypothesis is that we enter a high-standards context in which the assertion 'Zula knows that what she is looking at is a zebra' is no longer true. On my view, in contrast, we can explain what is going on without any appeal to the context-sensitivity of 'knows'. Instead, all that is being appealed to is the fact that becoming aware of an error-possibility can raise the evidential burden required in order to know, albeit in such a way that this burden is usually very easily met. But this fact no more indicates that 'knows' is a context-sensitive term than the fact that becoming aware of a defeater to something that one believes raises the evidential burden for knowledge. I am grateful to an anonymous referee for *Noûs* for pressing me on this issue.

²⁰ Note that I am not suggesting here that all versions of the transmission principle are acceptable (Wright is, of course, one of the most prominent defenders of the idea that the transmission principle, in at least some of its guises, should be rejected). In fact, I agree with Wright that there is a version of the transmission principle which should be rejected. What I argue here, however, is merely that zebra-style cases pose no problem for a specific formulation of the transmission principle (i.e., evidential transmission). The wider question of the status of the transmission principle more generally will need to be left to another occasion.

²¹ In Dretske's original discussion of this example, for instance, he directly moves to the issue of whether Zula can legitimately claim to know that what she is looking at is not a cleverly disguised mule. That it would be inappropriate for Zula to (flatly) make a claim of this sort does not indicate that such a claim would be false, however, especially given the point just mentioned that flat-out claims to know typically imply that the asserter has discriminatory epistemic support in

favour of what she asserts. I discuss the relationship between appropriate claims to know and discriminatory epistemic support in more detail in Pritchard (2007b; 2007c).

²² This is not to say, of course, that there is no scope for extending our treatment of the zebra-style cases to the sceptical problem. The point is rather that any such extension of the strategy would require further argumentation. As it happens, I think this distinction between favouring and discriminating epistemic support has important ramifications for the sceptical debate. After all, all parties to the sceptical debate surely wish to endorse the following two claims: (i) that one's knowledge that one has hands and that one is not a BIV is evidentially grounded; and (ii) that one cannot discriminate between having hands and being a BIV who merely thinks that he has hands. On the face of it, these two claims are in conflict with one another, but this is where the distinction between favouring and discriminating epistemic support can help to resolve this tension. For supposing that one could offer the further argumentation required to explain why (i) is true, the trick would clearly be to do this in such a way as to appeal only to favouring epistemic support and not discriminating epistemic support. One would then be free to endorse both of these claims, since the truth of (i) would no longer imply the falsity of (ii). I demonstrate the relevance of the distinction between favouring and discriminating epistemic support to the sceptical debate in Pritchard (2007c). In particular, I show that such a distinction is vital to anyone who offers a factive conception of epistemic support, since on this picture while one gets a straightforward explanation of why (i) is true, one does so at the expense of seemingly being unable to explain why (ii) is also true. In effect, then, this distinction enables one to endorse a view which many would regard as a form of infallibilism while at the same time being able to respond to the sceptical problem. *Contra* the conventional wisdom, then, infallibilism does not entail scepticism. (See Cohen (1988, 91) for an example of someone who holds both that allowing factive epistemic support is tantamount to infallibilism and that infallibilism entails scepticism). I am grateful to an anonymous referee for *Noûs* for pressing me on this issue.

²³ As noted above in endnote 17, a complicating factor in this regard is the existence of agents who ought to be aware of certain error-possibilities but aren't, perhaps because of deliberate close-mindedness on their part. Clearly, we need to extend the 'awareness' condition to include such agents, since intellectual traits like close-mindedness should not be able to insulate the positive epistemic status of one's beliefs. The issue of how best to extend the conception of awareness in play in order to capture these cases is, of course, complex, and so it would take us too far afield to explore it further here. I am grateful to an anonymous referee for *Noûs* for pressing me on this issue.

²⁴ One potential objection to the view defended here—put forward by an anonymous referee for *Noûs*—is that it makes knowledge 'unstable', in the sense that it can be easily lost. That the relevant favouring evidence is usually easily available, however, explains why this is not in fact the case. For while it is true that this view allows that a subject can have perceptual knowledge even while lacking the relevant favouring evidence, and hence concedes that such an agent's knowledge will be unstable in the sense that it is easily lost once the error-possibility in question is raised, such cases are bound to be relatively rare. After all, it is central to the view defended here that it takes very little by way of intellectual sophistication to be able to possess the relevant favouring evidence and bring it to bear on the error-possibility in hand, and so we can reasonably expect that in a wide range of cases one's knowledge will exhibit no such 'instability'. Moreover, that knowledge should be unstable in these special cases is surely far from counterintuitive. Note too that the view of knowledge in play here is entirely compatible with the idea that all knowledge (even the 'unstable' knowledge just described) is stable in the sense of being *safe*—i.e., it involves a belief in the target proposition that could not have easily been false.

²⁵ Such an assumption is common currency in the literature, especially when it comes to attributer contextualist or subject-sensitive (/interest-relative) invariantist accounts. For the main versions of attributer contextualist theories, see DeRose (1995), Lewis (1996), and Cohen (1988; 2000). For the main versions of subject-sensitive (interest-relative) invariantist views, see Hawthorne (2004) and Stanley (2005).

²⁶ Earlier versions of this paper were given at the University of St. Andrews and the University of Cambridge in 2006, and at the Universities of Copenhagen, Edinburgh and Geneva in 2007. I am grateful to the audiences on these occasions, especially Brit Brogaard, Matthew Chrisman, Andy Clark, Julien Dutant, Pascal Engel, Patrick Greenough, Lars Bo Gundersen, Jesper Kallestrup, Klemens Kappel, Katherine Hawley, Carrie Jenkins, Martin Kusch, Peter Lipton, Ram Neta, Daniel Nolan, Nikolaj Nottelmann, Matt Nudds, Erik Olsson, Mike Ridge, Joe Salerno, Jonathan Schaffer and Crispin Wright. Thanks also to Adam Carter, Adrian Haddock, Christoph Kelp, Ram Neta and Alan Millar for helpful discussion on themes relevant to the topic of this paper, and to two anonymous referees for *Noûs*. This research was supported by the award of an AHRC research leave grant.