

Perceptual Knowledge, Discrimination, and Closure

Santiago Echeverri
New York University, *Visiting Scholar*
Swiss National Science Foundation

Abstract: Carter and Pritchard (2016) and Pritchard (2010, 2012, 2016) have tried to reconcile the intuition that perceptual knowledge requires only limited discriminatory abilities with the closure principle. To this end, they have introduced two theoretical innovations: a contrast between two ways of introducing error-possibilities and a distinction between discriminating and favoring evidence. I argue that their solution faces the “sufficiency problem”: it is unclear whether the evidence that is normally available to adult humans is sufficient to *retain* knowledge of the entailing proposition and *come to know* the entailed proposition. I submit that, on either infallibilist or fallibilist views of evidence, Carter and Pritchard have set the bar for deductive knowledge too low. At the end, I offer an alternative solution. I suggest that the knowledge-retention condition of the closure principle is not satisfied in zebra-like scenarios.

Keywords: Perceptual knowledge; knowledge closure; epistemology of perception; relevant alternatives; perceptual discrimination

Consider a version of Fred Dretske’s (1970) famous zebra scenario:

The zebra scenario

Julia is in the zoo and gets a good look at a zebra in a pen clearly marked ‘zebra’.

Julia has normal vision, lacks expertise in zoology, and has performed no special checks on the animal. Given that the observation conditions are normal, Julia comes to know that that animal is a zebra.

Carter and Pritchard (2016) and Pritchard (2010, 2012, 2016) have used this scenario to formulate two seemingly incompatible ideas: the intuition that Julia’s perceptual knowledge of the zebra requires only limited discriminatory abilities and the intuition that Julia can add to what she knows by competent deduction.

The first intuition is at the heart of a family of views known as ‘relevant alternatives accounts of knowledge’. Carter and Pritchard spell out that intuition as follows: “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*” (Carter and Pritchard 2016:

970). A proposition, q , is an alternative to p just in case q is a logical contrary of p (Vogel 1999: 155). Thus, in order to know p , it is only required that one can rule out a proper subset of the logical contraries of p . Let us call this the ‘relevant alternatives intuition’.¹

In the case of perception, there is a natural way of fleshing out the relevant alternatives intuition: to be able to rule out an alternative is to be able to make perceptual discriminations between the target object and the object at issue in relevant alternatives (Carter and Pritchard 2016: 970). Thus, perceptual knowledge of zebras requires only the ability to discriminate zebras from the objects at issue in relevant alternatives. Carter and Pritchard think of relevant alternatives as possibilities that could plausibly obtain in the subject’s environment (for related ideas, see Austin 1946; Cohen 1988, 1999; Dretske 1970, 1971, 1981; Goldman 1976; Lewis 1996; Schaffer 2005). Assuming that Julia’s environment is normal, her perceptual knowledge that that animal is a zebra requires only that she can discriminate zebras from other animals that could plausibly be present in her environment (e.g., elephants, camels, dogs, and so on). Since perceptual knowledge requires only mundane discriminatory abilities, it is easy to come by.

Unfortunately, the relevant alternatives intuition seems to be in tension with another intuition: that Julia can add to what she knows by competent deduction. The closure principle spells out that intuition:

The closure principle

For all propositions, p , q , and subjects S : If S knows p , and S competently deduces q from p (thereby coming to believe q while retaining her knowledge of p), then S knows q .²

¹ The relevant alternatives intuition is shared by many philosophers who otherwise disagree about the relation between evidence and knowledge or the factors that determine whether an alternative is relevant. Taking a stance on these issues is certainly important to offer an account of knowledge. Fortunately, none of our arguments depends on how one fills in these details. See Cohen (1988, 1999), Dretske (1970, 1981), Lewis (1996), and Stine (1976) for different approaches. See Vogel (1999) for criticism.

Some people read Cohen as defending a non-relevant alternatives view. Vogel (1999) terms his view ‘plain contextualism’ and Pryor (2001: 98) strength-of-evidence contextualism. On this view, even if an alternative is irrelevant, the subject should have *some* evidence against alternatives to what he believes. But this reading is in tension with Cohen’s insistence that we lack evidence against radical skeptical hypotheses.

Suppose that Julia has normal logical skills. Hence, she can perform competent deductions from what she knows. If she knows that that animal is a zebra, she could reason that that animal is not a mule cleverly disguised to look like a zebra. This raises a problem, though: the closure principle seems to entail that Julia can have knowledge that the relevant alternatives intuition insists she lacks (Carter and Pritchard 2016: 972). Here is why. Closure seems to imply that Julia *must* be able to rule out *every* alternative to the zebra proposition, including the cleverly disguised mule hypothesis. By contrast, the relevant alternatives intuition tells us that Julia *does not need* to be able to rule out *every* alternative to the zebra proposition, including the cleverly disguised mule hypothesis. Interestingly, this problem can be formulated on *any* view of evidence. After all, there can be different ways of fleshing out the concept of ruling out an alternative.

Let us call ‘closure problem’ the apparent tension between the relevant alternatives intuition and closure. There have been different ways of solving this tension. Dretske (1970, 1981, 2005) famously argued that we should reject closure. Alas, many philosophers find this solution unpalatable. A popular view holds that closure offers a simple and plausible explanation of why deduction is a good way of extending knowledge in fields like mathematics and geometry (Hawthorne 2004; Williamson 2000). If the closure principle was rejected, we would be deprived of that simple and plausible explanation. That is why other philosophers have either rejected the relevant alternatives approach (Vogel 1990, 1999) or have tried to reconcile its underlying intuition with closure (Cohen 1988, 1999; Lewis 1996; Stine 1976).

In a series of insightful writings, Carter and Pritchard (2016) and Pritchard (2010, 2012, 2016) have proposed a novel solution to the closure problem. They have argued that—appearances to the contrary notwithstanding—competent deduction can extend our knowledge beyond the limits of our discriminatory powers. Carter and Pritchard’s solution relies on two theoretical innovations: a contrast between two ways of introducing error-possibilities and a distinction between discriminating and favoring evidence. On Carter and Pritchard’s view, the closure problem originates from a restrictive view of the evidence

² See Hawthorne (2004: 34), Pritchard (2012: 68), and Williamson (2000: 117).

available to normal adult humans. The restrictive view holds that normal adult humans—at least those who are capable of extending their knowledge by deduction—only have discriminating evidence at their disposal.³

I shall argue that Carter and Pritchard's view faces the "sufficiency problem": it is unclear whether the evidence that is normally available to adult humans is sufficient to *retain* knowledge of the zebra proposition and *come to know* the denial of the cleverly disguised mule hypothesis. I will show that, on either infallibilist or fallibilist views of evidence, Carter and Pritchard have set the bar for deductive knowledge too low.

In a more positive vein, I will suggest that the closure problem originates from a failure to distinguish the plausible claim that closure is valid from the less plausible claim that closure is non-vacuously satisfied in zebra-like scenarios. If we can show that a necessary condition of closure is violated in zebra-like scenarios, we will be able to preserve both the closure principle and the intuition that Julia cannot come to know the denial of the cleverly disguised mule hypothesis in zebra-like scenarios.

The article falls into seven sections. Section 1 summarizes Carter and Pritchard's solution. Section 2 spells out two assumptions in that solution. Section 3 introduces the sufficiency problem. In sections 4 and 5, I defend a pessimistic answer to the sufficiency problem. Section 6 responds to objections and Section 7 sketches an alternative account.

1. Carter and Pritchard's Solution

Carter and Pritchard's solution to the closure problem can be summarized in four steps.

Step 1: When a subject considers some entailments, she often becomes aware of error-possibilities. Thus, when Julia thinks: 'If that animal is a zebra, that animal is not a mule cleverly disguised to look like a zebra', she becomes aware of the error-possibility that she is looking at a mule cleverly disguised to look like a zebra. If Julia understands the

³ This assumption is relatively common in the literature. Examples include Schaffer (2005), Stine (1976), and Wright (2003). Dretske (1970: 39) has a more expansive conception of evidence but also thinks that it is not sufficient to know the denial of the cleverly disguised mule hypothesis. The present paper may be seen as an attempt to vindicate Dretske's intuition without either rejecting the closure principle or relying on Dretske's controversial theory of conclusive reasons.

propositions involved, she will be led to recognize the error-possibility *as* incompatible with what she believes (Carter and Pritchard 2016: 982; Pritchard 2016: 155).

Step 2: When a subject becomes aware of an error-possibility *as* incompatible with what she believes, she momentarily loses her initial knowledge. Therefore, she ought to form a view on that possibility if she is to retain her knowledge of the target proposition (Pritchard 2012: 78). Two considerations lend support to this verdict. First, the possession of knowledge (by rational subjects, anyway) is incompatible with a subject's awareness of undefeated defeaters, even if those defeaters turn out to be misleading (Pritchard 2010: 265 n 19; 2012: 51, 99, 89, 104 n 17, 105 n 21, 146). Second, an undefeated defeater may lead one to lose one's belief in the target proposition. If knowledge entails belief, an undefeated defeater may lead one to lose one's knowledge of the target proposition too (Pritchard 2016: 212 n 4).⁴

Step 3: To retain her initial knowledge, the subject must rationally dismiss the error-possibility. If the error-possibility targets the subject's perceptual evidence, she should not rely on that perceptual evidence to dismiss the error-possibility. The perceptual evidence is, so to speak, 'neutralized' by the error-possibility (Pritchard 2012: 85, 99). A subject can rationally dismiss an error-possibility in two different ways. If the error-possibility is rationally motivated, the subject should provide independent evidence that speaks to the reasons offered in support of the error-possibility. If the error-possibility is not rationally motivated (but is 'merely raised'), the subject may ignore the error-possibility because it lacks rational support. Think of hypotheses formulated 'on a whim'. It seems wrong to expect someone to provide independent evidence to dismiss whimsical hypotheses.

Step 4: Adult humans normally have background evidence that enables them to dismiss rationally motivated error-possibilities. Focusing on the cleverly disguised mule hypothesis, Pritchard offers some plausible examples:

This will include evidence regarding the plausibility of this deception, the likelihood of it being spotted, the penalties that would be imposed were it to be discovered, the lengths

⁴ But see Radford (1966), for a controversial counterexample to the claim that knowledge that *p* entails belief that *p*.

that one would need to go in order to make this deception plausible (including the costs that one would incur), and so on (Pritchard 2016: 131-2; see also Carter and Pritchard 2016: 977; Pritchard 2012: 79ff.).

If we take seriously this background evidence, the tension between the relevant alternatives intuition and closure seems to evaporate. For, Carter and Pritchard think, this tension arises from a very restrictive conception of the evidence available to normal adult humans. The following principle articulates this restrictive conception:

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 (Pritchard 2010: 251; see also Pritchard 2012: 73).

If the evidential position of normal adult humans includes background evidence, the discrimination principle is false: a subject's knowledge of the denial of the error-possibility can be supported by her background evidence. This background evidence will enable the subject to *retain* her initial knowledge and *dismiss* the error-possibility. In a nutshell: the tension between the relevant alternatives intuition and closure arises from the unsupported assumption that subjects only have 'discriminating evidence', neglecting the possibility that they can come to know the logical consequences of their knowledge via their background evidence (or 'favoring evidence').

To summarize, suppose that a subject knows p (e.g., that animal is a zebra) and knows that p entails *not- q* (e.g., that animal is not a mule cleverly disguised to look like a zebra). If q is a merely raised error-possibility, realizing that q was not rationally motivated will enable our subject to know *not- q* by competent deduction from her knowledge of p . If q is rationally motivated, our subject will momentarily lose her knowledge of p . Yet, she will be able to recover that knowledge and come to know *not- q* by competent deduction *if* she has independent, background evidence that speaks to the specific grounds offered in

support of *q* (Carter and Pritchard 2016: 977-8). On the plausible assumption that normal adult humans have that background evidence, the closure problem is solved.

2. Two Assumptions

Carter and Pritchard's solution hinges on two assumptions. It is worth making those assumptions explicit before we evaluate their view.

The first assumption is that a normal adult human can come to know, via competent deduction, propositions that outstrip her discriminatory powers. It is tempting to think that this assumption goes hand in hand with *any* defense of closure. But that is a mistake. On the one hand, there have been defenses of closure that predict that a normal adult human cannot come to know propositions that outstrip her discriminatory powers. Examples include contrastivism (e.g., Schaffer 2007) and—some readings of—attributer contextualism (e.g., Cohen 1988; Lewis 1996). On the other hand, it might turn out that a normal adult human's background evidence is insufficient to *retain* her initial knowledge and *come to know* the relevant error-possibility. Let us call 'Optimism' the assumption that a normal adult human's background evidence is sufficient to retain her initial knowledge and come to know the target error-possibility. Carter and Pritchard subscribe to Optimism.⁵

The second assumption is that one can solve the closure problem without taking on any substantial epistemological commitments:

[W]e would ideally seek a resolution of the closure problem which has broad appeal, *whatever one's wider epistemological inclinations, and which is thus not allied to any particular contentious proposal in epistemology* (Carter and Pritchard 2016: 975; emphasis mine).

⁵ See Carter and Pritchard (2016: 977-8, 980), Pritchard (2012: 68ff., 76, 80-1, 86-7, 95-8, 99-100, 102 n 6, 104 n 17, 105 n 19), and Pritchard (2016: 131-2, 137). The qualification 'normal adult human' is meant to exclude infants and animals. There is some lack of clarity on how inclusive the class of normal adult humans is. Carter and Pritchard recognize that adult humans may *sometimes* be unable to know, via competent deduction, the logical consequences of their perceptual knowledge. They make room for cases involving adult humans who lack *any* background evidence (Pritchard 2012: 79, 83) or whose background evidence is defeated by the epistemic authority of their interlocutor (Carter and Pritchard 2016: 983). Nevertheless, their discussion is guided by the assumption that the evidence available to normal adult humans will *often* suffice to retain their initial knowledge and dismiss the target error-possibility. That assumption is the main target of this paper.

This methodological ideal has some appeal. The closure principle spells out a central intuition shared by different epistemological theories. If a solution to the closure problem relied on a particular epistemological theory, it would not be widely accepted by the philosophical community.

It might be objected that this methodological ideal is overly optimistic. Indeed, it might turn out that not everyone will be able to embrace a solution to the closure problem without incurring some substantial epistemological commitments. Nevertheless, there is a restricted version of the methodological ideal that is harder to dismiss: maybe one could solve the closure problem without taking a stance on the nature of evidence. That is Carter and Pritchard's view. They think that they can solve the closure problem without committing themselves to any specific view of the relation between evidence and knowledge. Let us term this assumption 'Ecumenism'.⁶

We shall see that Optimism cannot be adequately defended while endorsing Ecumenism. More interestingly, Optimism is unjustified on either infallibilist or fallibilist theories of evidence.

3. The Sufficiency Problem

In this section, I present the "sufficiency problem" for Optimism.

Let us start with some relatively uncontroversial assumptions. Given the construction of the zebra scenario, we can credit Julia with perceptual knowledge that that animal is a zebra. If Julia's knowledge is evidentially based, we may also assume that her total evidence (plus one's preferred conditions for knowledge) is sufficient to know that that animal is a zebra. We can also grant this much: Julia has overall better evidence in support of her belief that that animal is a zebra than for believing that that animal is a mule cleverly disguised to look like a zebra. This seems to follow from Julia's possession of background evidence against the possibility that that animal is a mule cleverly disguised to look like a zebra.

⁶ Ecumenism plays a key role in Pritchard's defense of epistemological disjunctivism. He needs to defend closure without presupposing the truth of the latter (Pritchard 2012: 63-4, 80, 82).

Let us suppose now that Julia tries to deduce the conclusion that that animal is not a mule cleverly disguised to look like a zebra from her perceptual knowledge of the zebra. The sufficiency problem is this: Can Julia *retain* her knowledge of the zebra proposition and *come to know* the denial of the cleverly disguised mule hypothesis with the aid of the evidence at her disposal? Optimism answers “yes”; Pessimism answers “no”. I will argue that Pessimism should be taken seriously.

Our starting point is a simple logical remark. One’s evidence in support of one’s belief in p can be *better* than one’s evidence for believing *not- p* even though that evidence is not *good enough* for knowledge of *not- p* . In other words, one’s total evidence can be sufficient for preferring p to the alternative *not- p* without being sufficient for knowledge of *not- p* . Therefore, even if p enjoys overall better evidential support than *not- p* , it does not follow that we know *not- p* . Pritchard (2016: 47-8, 57-9, 162-3) has made a similar point in his argument for the claim that the closure principle is logically stronger than the underdetermination principle. Our point is rather different. Even if we grant Pritchard’s (2012: 76) contention that “[Julia’s] evidential position is in fact much stronger than we ordinarily tend to think”, it need not be sufficient to *retain* her knowledge of the zebra proposition and *come to know* the denial of the cleverly disguised mule hypothesis.

Some readers might find our logical remark puzzling. We have granted that Julia’s total evidence is good enough for knowledge that that animal is a zebra. So, why couldn’t Julia know the denial of the cleverly disguised mule hypothesis on the same evidential basis? Part of the answer lies in some plausible assumptions. Many philosophers hold that knowledge that q entails belief that q . The literature on propositional justification offers many examples of cases in which a subject has evidence for q but does not believe q (Firth 1978). If those situations are possible, then the possession of knowledge-level evidence for a proposition is not sufficient for knowledge of that proposition. Julia might have knowledge-level evidence in favor of the propositions <That animal is a zebra> and <That animal is not a mule cleverly disguised to look like a zebra> but form no opinion about the latter hypothesis. Alternatively, Julia could have knowledge-level evidence in favor of the proposition <That animal is not a mule cleverly disguised to look like a zebra> but base

her belief in that proposition on bad reasons (Schaffer 2007: 235-6). In sum, one can know p , know that p entails q , but not know q because one has no belief in q .

Suppose now that Julia tries to deduce the denial of the cleverly disguised mule hypothesis from her perceptual knowledge of the zebra. Pessimism is the claim that Julia—and any normal adult human—can neither *retain* her knowledge of the zebra proposition, nor *come to know* the denial of the cleverly disguised mule hypothesis with the evidence at her disposal. A reason why Optimism may seem plausible is that we tend to discount the epistemic effects of engaging in a process of deduction. Now, it might turn out that these epistemic effects prevent us from retaining knowledge of the entailed proposition and come to know the entailed proposition, even if the evidence at our disposal is good enough to know those propositions *outside the process of deduction*. A merit of Carter and Pritchard's defense of closure is that it underscores the epistemic effects of the process of deduction. So, it is ideally suited to examine the prospects of Pessimism.

In the remainder of this paper, we will see that, even if one's total evidence is good enough to know that that animal is a zebra and also favors the zebra proposition over the cleverly disguised mule hypothesis, that evidence is not sufficient to *retain* knowledge of the zebra proposition and *come to know* the denial of the cleverly disguised mule hypothesis.

4. Infallibilism

Infallibilism holds that we can only have knowledge of p on the basis of evidence that entails p (Williamson 2000). This section defends a conditional claim: If infallibilism is true, then Optimism is false.

Suppose that Julia knows that that animal is a zebra on the basis of factive perceptual evidence. Suppose also that Julia tries to deduce from her perceptual knowledge the proposition that that animal is not a mule cleverly disguised to look like a zebra. In doing so, Julia becomes aware of the error-possibility that that animal is a mule cleverly disguised to look like a zebra. According to Carter and Pritchard, this undefeated defeater is sufficient for Julia to momentarily lose her knowledge that that animal is a zebra (Section 1, Step 2). Moreover, it also neutralizes Julia's perceptual evidence for the proposition that

that animal is a zebra, *even if that evidence is factive* (Section 1, Step 3). As Pritchard (2016: 139) aptly puts it, “so long as one lacks an independent rational basis to dismiss the rationally motivated error-possibility in question, then one ought to no longer regard oneself as being in possession of the target factive rational support for one’s perceptual belief”. Therefore, Julia needs to invoke independent evidence to dismiss the error-possibility. Suppose now that the error-possibility has been rationally motivated. Julia could invoke some background evidence that speaks to the specific reasons that motivate the error-possibility. Unfortunately, the background evidence that Julia could plausibly cite to dismiss that error-possibility is too weak by infallibilist lights. The considerations ‘it is implausible that I am being deceived now’, ‘it is likely that, if that is a cleverly disguised mule, the deception will have been spotted’, ‘it would be very expensive to make this deception plausible’, and so on, do not *entail* that that animal is a zebra (and thus that that animal is not a mule cleverly disguised to look like a zebra). Therefore, given an infallibilist theory of evidence, Optimism is false: Julia would irremediably lose her perceptual knowledge that that animal is a zebra and could not come to know the denial of the error-possibility.

It might be argued that the previous argument does not go through if we endorse epistemological disjunctivism. Epistemological disjunctivism holds that, in the good cases, subjects have perceptual knowledge that *p* in virtue of having perceptual evidence that is both factive and reflectively accessible (Pritchard 2012: 13). Interestingly, Carter and Pritchard (2016: 975) sometimes flirt with the idea that, if epistemological disjunctivism is true, then the subject’s total evidence would suffice to know the denial of the cleverly disguised mule hypothesis.⁷ Thus, they might want to resist our previous argument by appealing to epistemological disjunctivism. Although Julia’s background evidence does not entail that that animal is a zebra, her perceptual evidence does entail it. When Julia gets a good look at the zebra, she sees that that animal is a zebra. Given that seeing that *p* is factive, the conjunction of that perceptual evidence with Julia’s background evidence entails that that animal is a zebra (and thus that that animal is not a mule cleverly disguised to look like a zebra).

⁷ Pritchard formulates epistemological disjunctivism in terms of reasons, which may not be equivalent to evidence. This difference does not matter in the current context.

Unfortunately, this objection overlooks some of the epistemic effects of the recognition of error-possibilities *as* incompatible with what one believes. The first effect is that the subject momentarily loses her initial perceptual knowledge (Section 1, Step 2). The second effect is that the subject cannot regard herself as seeing that *p*. The reason is straightforward: if the error-possibility targets the subject's perceptual evidence, she cannot rely on that perceptual evidence to dismiss the error-possibility. The perceptual evidence is, so to speak, 'neutralized' by the error-possibility (Pritchard 2012: 85, 99; *cf.* Section 1, Step 3).⁸ That is why *independent* evidence is required to dismiss rationally motivated error-possibilities. Crucially, it does not really matter whether the perceptual evidence is factive or not: "so long as one lacks an independent rational basis to dismiss the rationally motivated error possibility in question, *then one ought to no longer regard oneself as being in possession of the target factive rational support for one's perceptual belief*" (Pritchard 2016: 139; emphasis mine). If the subject ought not to regard herself as seeing that that animal is a zebra, her only way of retaining her knowledge that that animal is a zebra is to rely on the remaining, background evidence at her disposal. But *that* background evidence does not entail that that animal is a zebra. So, by infallibilist lights, Optimism is false.

To be sure, Carter and Pritchard might reply that their solution does not need to be compatible with infallibilist theories of evidence. In the next section, I argue that their solution is also implausible within a fallibilist theory of evidence.

5. Fallibilism

Carter and Pritchard's solution is incompatible with infallibilism about evidence, even if one construes perceptual evidence along epistemological-disjunctivist lines. The remaining alternative is *fallibilism*. According to fallibilism, one can know *p* on the basis of evidence that does not entail *p* (Cohen 1988, 1999; Pryor 2000; Vogel 1999). A popular version of this idea says that knowledge of *p* requires that the epistemic probability of *p* be above a

⁸ If seeing that *p* entails knowing that *p*, then one cannot retain one's seeing that *p* without knowing that *p* (Williamson 2000). Pritchard (2012) rejects the entailment. I am granting Pritchard's controversial view for the sake of the argument.

certain threshold t (where t is less than 1). If Carter and Pritchard want to defend Optimism within a fallibilist framework, they must show that the subject's evidence adds up to a threshold above t .

Cohen (1988: 103) offers an influential solution to the threshold problem. On his view, it is not possible to specify what constitutes sufficient evidence to know the denial of an alternative because that will depend on the context in which the knowledge attribution occurs. Is this solution available to Carter and Pritchard? An initial problem is that Pritchard (2012: 104 n 17; 2016: 36-46, 144-52) has explicitly rejected attributer contextualism.

Carter and Pritchard might try to implement variable fallibilist evidential standards without presupposing a contextualist semantics for 'know'. Indeed, one can read Carter and Pritchard's account as introducing a distinction between two different evidential standards for knowledge. In the 'default' cases, we can easily come to know that that animal is a zebra based only on our limited, discriminating evidence. This is the core of the relevant alternatives intuition. By contrast, an adult human who becomes aware of the cleverly disguised mule hypothesis *as* incompatible with the zebra proposition needs independent, background evidence to retain her perceptual knowledge.

I want to suggest that the background evidence that subjects could plausibly have is not sufficient for knowledge *after they have become aware of the error-possibility*. The argument exploits the following, intuitive principle:

Coherent self-ascription principle

If a subject, S , can retain her knowledge of p after she has recognized an error-possibility *as* incompatible with p , then S can self-ascribe knowledge of p without manifest incoherence.⁹

This principle seems extremely plausible. Moreover, it provides a simple explanation of a common line of thought.

⁹ Hereafter, I abbreviate 'recognize *as*...' to 'recognize'.

Williams (1991: 352-5) observes that knowledge is very unstable during epistemological reflection; doing epistemology makes it very easy to lose one's knowledge. Cohen (1988) concurs. On his view, the salience of an error-possibility often suffices to lose one's knowledge. Lewis' (1996: 434ff.) Rule of Attention makes the same point: paying attention to a far-fetched error-possibility often suffices to lose one's knowledge. The coherent self-ascription principle offers a simple explanation of these remarks. Reflection, salience or attention often lead us to lose our knowledge because, in those circumstances, the self-ascription of knowledge often becomes manifestly incoherent.

What is the source of incoherence? When someone takes herself to know *p* but also recognizes that she might be mistaken, she is aware of the factivity of knowledge and the fallibility of her evidence. Alas, these two thoughts suffice to make any knowledge self-ascription manifestly incoherent. As Lewis' (1996: 419) aptly puts it: "To speak of fallible knowledge, of knowledge despite uneliminated possibilities of error, just *sounds* contradictory". If a subject were to think: 'I know, yet I have not eliminated all possibilities of error', her knowledge self-ascription would be manifestly incoherent. That is probably why Williams, Cohen, and Lewis conclude that the consideration of error-possibilities often leads us to lose one's knowledge.

This conclusion seems appealing, *even by fallibilist lights*. Someone who entertains those types of thoughts is aware of a manifest incoherence. If the coherent self-ascription principle is true, entertaining those thoughts will also prevent our subject from retaining knowledge. In what follows, I shall defend a *restricted version* of Lewis' claim: the self-ascription of knowledge becomes manifestly incoherent when the subject becomes aware that her background evidence has not eliminated *all* error-possibilities. My suggestion is that this is a normal consequence of trying to deduce the denial of the cleverly disguised mule hypothesis in zebra-like scenarios.

Suppose that Julia's interlocutor said: 'I have heard from John that, for budgetary reasons, the zebra has been replaced with a painted mule'. Julia might point out that John has often lied in the past. Hence, Julia would have *some* reason to put the error-possibility to one side. Yet, it would be odd for Julia to hold firm on the following, factive self-ascriptions: 'I *know* that that animal is a zebra' or 'I *know* that that animal is not a mule

cleverly disguised to look like a zebra'. After all, even the worst liars sometimes tell the truth. Being a normal adult, Julia would surely be aware of this fact. Since Julia would be aware that her background evidence does not rule out the possibility that this time John is telling the truth, she could not coherently stick to those factive self-ascriptions. If Julia understands what it takes to self-ascribe a factive state like knowledge, she should realize that knowing p excludes *any* possibility, however farfetched, that *not- p* (or the incompatible proposition q). Recall now the coherent self-ascription principle: a manifest incoherence prevents one from retaining knowledge of p after one has recognized an error-possibility. Given this principle, Julia could not retain knowledge that that animal is a zebra.

This result generalizes to all the other pieces of background evidence that normal adult humans plausibly have in zebra-like scenarios. The considerations 'it is implausible that I am being deceived now', 'it is likely that, if that is a cleverly disguised mule, the deception will have been spotted', 'it would be very expensive to make this deception plausible', and so on, all yield the same verdict: normal adult humans are surely aware that, given that background evidence, it is still possible that that animal is a mule cleverly disguised to look like a zebra. But 'I know that that animal is a zebra, yet I have not eliminated all possibilities of error' seems contradictory. Therefore, normal adult humans could not retain their knowledge that that animal is a zebra after they have recognized the error-possibility that that animal is a mule cleverly disguised to look like a zebra.

Of course, Julia might still make *some* factive self-ascriptions: 'I know that it is *very likely* that that animal is a zebra' or 'I know that it is *very unlikely* that that animal is a mule cleverly disguised to look like a zebra'. Alas, the contents of these knowledge self-ascriptions are different from the original everyday proposition (that animal is a zebra) and the denial of the original error-possibility (that animal is a mule cleverly disguised to look like a zebra).

Therefore, even by fallibilist lights, Optimism seems false. Even if fallible evidence may be sufficient to gain perceptual knowledge of zebras in the 'default' cases, it need not be sufficient to *retain* that knowledge and *come to know* the denial of the cleverly disguised mule hypothesis in the process of deduction. If the reader disagrees, she can see our argument as an invitation to address the sufficiency problem.

6. Objections and Replies

Let us consider two objections.

Objection 1: The coherent self-ascription principle smuggles an infallibilist theory of knowledge into the fallibilist framework. Moreover, it has the unwanted consequence that Julia cannot self-ascribe knowledge of the zebra proposition.¹⁰

Reply: As I defined it, infallibilism holds that we can only have knowledge of p on the basis of evidence that entails p (Section 4). The coherent self-ascription principle does not rely on that claim. It exploits less controversial ideas: the widespread idea that knowledge is factive and our awareness, in the context of deduction, of the fallibility of our background evidence. The coherent self-ascription principle offers a simple explanation of why so many philosophers—including fallibilists—think that reflection, salience or attention are sufficient to lose knowledge of the zebra proposition and prevent us from coming to know the denial of the cleverly disguised mule hypothesis. This claim is consistent with fallibilism: if you are completely unaware of the cleverly disguised mule hypothesis, there is no manifest incoherence in the use of perception to come to believe that that animal is a zebra. Assuming that other conditions for knowledge are satisfied, nothing prevents you from coming to know that that animal is a zebra.

Objection 2: Consider the following argument:

Premise 1. If the coherent self-ascription principle is true, then favoring support could never entitle one to rationally dismiss an error-possibility.

Premise 2. Favoring support could entitle one to rationally dismiss an error-possibility.

Conclusion. So, the coherent self-ascription principle is not true.¹¹

Reply: I am not sure how the two premises are justified. So, the following discussion is rather tentative.

¹⁰ I owe this objection to a referee for this journal.

¹¹ A referee for this journal urged me to discuss this argument.

Premise 2 seems to rely on an ambiguity concerning 'rationally dismiss'. On a weak reading, the premise seems obviously true, but the argument is invalid. On a strong reading, the argument is valid, but its premise is unjustified.

On a *weak* reading, *S* rationally dismisses a hypothesis, *b*, just in case *S* puts *b* to one side on the basis of reasons that speak against *b*. The coherent self-ascription principle is consistent with this weak reading. We can grant that our normal background evidence provides us with reasons to put the cleverly disguised mule hypothesis to one side. Still, the weak reading does not license the conclusion. The coherent self-ascription principle concerns our *knowledge* of the entailing proposition (and the entailed proposition) after we have recognized an error-possibility. Alas, one may put *b* to one side on the basis of reasons that speak against *b* without thereby knowing the falsity of *b*. This weak sense of 'rationally dismiss' is familiar from scientific practice. A scientist may rationally dismiss a rival hypothesis because it is not very fruitful, or it is unlikely to be true. Yet, the available scientific evidence may be insufficient to know the falsity of that hypothesis. My intuition is that, even if we do not know the falsity of a rival hypothesis, it is often rational to put it to one side in order to explore other hypotheses.

On a *strong* reading, *S* rationally dismisses a hypothesis, *b*, just in case *S* puts *b* to one side on the basis of reasons that speak against *b* and *S*'s reasons allow *S* to know the falsity of *b*. The argument is valid on this strong reading. Unfortunately, this interpretation of premise 2 is nothing but an unjustified assertion of what is at issue: whether the sort of background evidence that adult humans normally have is sufficient to retain knowledge of the entailing proposition and come to know the denial of the error-possibility after they have recognized that error-possibility.

If you find the coherent self-ascription principle unpalatable, you can see it as an invitation to say more on an important issue: When is the possession of evidence against a recognized error-possibility sufficient to retain one's initial knowledge and come to know the denial of that error-possibility? Simply asserting that one's background evidence could entitle one to know the falsity of an error-possibility does not address this problem.

Premise 1 makes a strong claim that I do not make. What I said is this: manifest incoherence in a potential knowledge self-ascription prevents one from retaining that

knowledge after the recognition of an error-possibility. So, there might be cases in which the coherent self-ascription principle is satisfied. My bet is that they will be more infrequent than Carter and Pritchard assume. A natural strategy is to construe cases in which one's background evidence is so strong by the subject's lights that it becomes extremely hard to think of any uneliminated error-possibilities. Consider an example adapted from Dretske (1971: 13). Suppose that I know Peter exceptionally well. He is my brother, we grew up together, and I have never caught him in a lie. I also know that Peter has an almost pathological devotion to telling the truth about the zoo where he has been working for the past fifteen years. Suppose now that Peter has told me several times that the zoo authorities would never replace a zebra with a cleverly disguised mule. In this case, it seems plausible to think that Peter's testimony offers a conclusive reason in Dretske's sense: Peter would not have said that that animal is a zebra unless it was a zebra. Given what I know about Peter, his testimony may prevent me from becoming aware of uneliminated error-possibilities. So, Peter's testimony could enable me to retain knowledge that that animal is a zebra and come to know that that animal is not a mule cleverly disguised to look like a zebra.

Many philosophers have the intuition that, in the zebra scenario, the only way of knowing the falsity of the cleverly disguised mule hypothesis is to approach the animal and check whether it is painted. Carter and Pritchard seem to think that those philosophers presuppose the truth of the discrimination principle. Here is an alternative explanation: even if favoring evidence is factored in, that favoring evidence does not seem to be good enough to come to know the falsity of the cleverly disguised mule hypothesis. My suggestion is that the pessimistic outlook follows either from infallibilism about evidence or from the coherent self-ascription principle. Given that Pessimism is consistent with closure, we do not need to be worried about Pessimism.

7. An Alternative Solution

Here is my solution to the closure problem: that problem arises from the unsupported assumption "that the closure principle seems to require that in order for [Julia] 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” (Carter and Pritchard 2016: 972; emphasis mine). If the activity of deduction inevitably leads us to lose our perceptual knowledge in zebra-like scenarios, then the closure principle makes no demand for knowledge beyond our discriminatory powers. Why is it so tempting to think that the closure principle makes that demand? I have suggested that we tend to discount the epistemic effects of engaging in a process of deduction (Section 3). This might be a side-effect of a tendency to conflate a synchronic formulation of closure (‘synchronic closure’) with the diachronic formulation of closure we have been working with (‘diachronic closure’). Crucially, only synchronic closure makes a demand for knowledge beyond our discriminatory abilities.

Dretske’s (1970) anti-closure view was originally formulated against synchronic closure:

The synchronic closure principle

For all propositions, p , q , and subjects S : If S knows p and knows that p entails q , then S knows q .¹²

Let us assume that Julia knows that that animal is a zebra and knows that that proposition entails that that animal is not a mule cleverly disguised to look like a zebra. Given those assumptions, synchronic closure makes a very strong demand on knowledge: Julia *must already* know that that animal is not a mule cleverly disguised to look like a zebra. That is why attacks on closure are very often attacks on Pessimism and defenses of closure are very often defenses of Optimism. Dretske’s basic strategy is to reject synchronic closure by rejecting Optimism. He gave many examples in which—he thought—a subject knows p , knows that p entails q , but does not know q .

It is widely admitted that synchronic closure is prey to obvious counterexamples.¹³ Moreover, there is an emerging consensus that diachronic closure can avoid many of those

¹² See Dretske (1970: 33). For similar formulations, see Cohen (1988: 93; 1999: 62), DeRose (1995: 31 n 33), Dretske (2005: 13), Nozick (1981: 204ff.), Stine (1976: 249), and Vogel (1990: 13; 1999: 156).

counterexamples. I shall not rehearse these considerations here (but see [Author 1]; Hawthorne 2004; Pritchard 2016; Schaffer 2007). Once we adopt the diachronic formulation, we can consistently endorse closure and reject Optimism. A subject can know p , know that p entails q , but be unable to *retain* her initial knowledge of p and *come to know* q when she tries to deduce q from p . Thus, we can steer a middle way between the anti-closure view and Optimism. Dretske was right when he argued that we cannot use deduction to come to know the denial of the cleverly disguised mule hypothesis in the zebra scenario. Yet, he was wrong to think that this invalidated (diachronic) closure.

To be sure, contextualists and contrastivists have also tried to preserve closure and reject Optimism (Section 2). We have achieved the same result without putting forward substantial theories about either the semantics of 'know' or the logical form of knowledge attributions. These linguistic theories are highly controversial. It has been argued that attributer contextualism about 'know' faces a problem of semantic blindness and requires an error-theory of semantic competence (Cohen 1999: 77; Conee 2005: 60; DeRose 1995: 40-1). Our solution does not face any of those problems.

Contextualists have often focused on synchronic closure. Thus, their strategy has been to explain how a subject can already know the denial of the cleverly disguised mule hypothesis *before* she performs any inference. The standards for knowledge are low before we perform any inference. So, we know the denial of the cleverly disguised mule hypothesis. The standards for knowledge are high when we perform some inferences. So, we do not know the denial of the cleverly disguised mule hypothesis. If we endorse diachronic closure, we do not need to assume that subjects *must* have any 'pre-deductive knowledge' of the denial of the cleverly disguised mule hypothesis. That strikes me as the right result.¹⁴

¹³ Some of them come from the distinction between propositional and doxastic justification. See Section 3.

¹⁴ Of course, one *could* have pre-deductive knowledge of the denial of the cleverly disguised mule hypothesis. What seems wrong is to construe that knowledge as a *necessary requirement* to have knowledge of the zebra proposition. Otherwise, we would have trouble accommodating the idea that deduction can provide new knowledge.

Carter and Pritchard advertise their view as ecumenical concerning the relation between evidence and knowledge. Yet, the sufficiency challenge compels them to take a stance on that controversial issue. Our view has the merit of being genuinely ecumenical. Whether you are attracted by either infallibilism or fallibilism, you can accept the present solution to the closure problem.

It has been pointed out that a denial of closure would lead us to reject intuitively plausible instances of correct reasoning (Feldman 1999; Hawthorne 2004; Schaffer 2007). These are cases in which the closure principle is not vacuously satisfied. Indeed, subjects can often add to what they know by competent deduction from what they know. Let us term 'closure-validating inferences' the cases in which the closure principle is not vacuously satisfied. This suggests a challenge for our view: If one holds that the closure principle is vacuously satisfied in zebra-like scenarios, one should explain why it is not vacuously satisfied in *all* cases.

Here is my response in a nutshell: the most plausible cases in which subjects can use deduction to know the consequences of what they know are cases in which subjects have not recognized error-possibilities *as* incompatible with the entailing proposition. This is particularly clear in inferences that exploit logical and conceptual connections. It is therefore no accident that such inferences have played a key role in most defenses of the closure principle.

If I know that there are cookies in the jar, I am in a position to know, via competent deduction, that there is something in the jar (Dretske 2005: 30). If I know that Anita has two brothers, I am in a position to know by competent deduction that Anita has at least one brother. If you know that Bill and Susan married each other, you can also know by competent deduction that Susan married someone (Hawthorne 2004: 33).

A commonly held view says that logic merely spells out ways of combining semantic contents; logical relations are formal because they hold no matter how the world is (MacFarlane 2000). If a subject masters the quantifiers, she tacitly knows that the conclusions of the previous inferences are somehow 'implicit' in the content of their premises. Hence, there is no reason why a subject who performs those deductions should recognize some error-possibilities *as* incompatible with the entailing propositions.

Other closure-validating inferences exploit conceptual connections. Red is a color. Thus, from one's knowledge that something is red, one can deduce and come to know that something is colored (Hawthorne 2004: 35). If I know that A is part of B and know that B is part of C, then I can know, via competent deduction, that A is part of C. If I know that Julia has a brother, I can know by competent deduction that her brother is male (Harman 1986: 17).

When a deductive transition exploits conceptual connections of these sorts, the subject is not led to recognize error-possibilities *as* incompatible with the entailing proposition. Intuitively, if the application conditions of the concept RED obtain, the application conditions of the concept COLOR obtain too. Whatever constitutes evidence for the wall being red is also evidence for the wall being colored. Hence, if a subject competently applies the concept RED based on visual experience, her awareness of the proposition that the wall is colored won't make her aware of any error-possibility. Interestingly, this consideration does not apply in the zebra scenario. Someone can apply the concept ZEBRA on the basis of limited discriminatory evidence. Thus, even if the application conditions of the concept ZEBRA obtain, the application conditions of the concept MULE CLEVERLY DISGUISED TO LOOK LIKE A ZEBRA do not need to obtain. Given our limited discriminatory powers, we are disposed to think that we could be wrong in believing as we do when we are confronted with the cleverly disguised mule hypothesis. This point holds even if our perceptual evidence is factive and even if, necessarily, zebras are not cleverly disguised mules.

To be sure, a complete defense of our solution would require an examination of other conditions necessary to retain knowledge in the process of deduction. I hope to have shown, however, that one can consistently retain three ideas: the intuition that deduction is not a good way of extending knowledge in zebra-like scenarios, the intuition that deduction is generally a good way of extending knowledge, and the intuition that the closure principle is valid.

8. Conclusion

It is important to distinguish the plausible claim that the closure principle is valid from Optimism, i.e. the less plausible claim that subjects can use deduction to know propositions that outstrip their discriminatory powers in zebra-like scenarios. I have argued that Carter and Pritchard's defense of Optimism fails. Normal adult humans cannot use their background evidence to retain their knowledge of the zebra proposition and come to know the denial of the cleverly disguised mule hypothesis when they try to perform the relevant deduction. This point holds on either infallibilist or fallibilist views of evidence. Therefore, the closure principle is vacuously satisfied in zebra-like scenarios. Nevertheless, we can still use deduction to extend knowledge in many other cases. This is particularly plausible when we do not recognize error-possibilities. Our view is neutral on the relation between evidence and knowledge, so it can be accepted by both infallibilists and fallibilists. In addition, it does not incur any theoretical commitments about the semantics and logical form of 'know'.

Acknowledgments

I would like to thank Arturs Logins, Davide Fassio, and two anonymous referees for their detailed and challenging comments on an earlier draft of this paper. This work was funded by a generous grant from the Swiss National Science Foundation (FNS P300P1_161061/1).

References

- Austin, J. L. (1946) Other Minds, *Proceedings of the Aristotelian Society, Supplementary Volume* 20: 148-87
- Carter, A. and D. Pritchard (2016) Perceptual Knowledge and Relevant Alternatives, *Philosophical Studies* 173 (4): 969-90
- Cohen, J. (1988) How to Be a Fallibilist, *Philosophical Perspectives* 2: 91-123
- Cohen, J. (1999) Contextualism, Skepticism, and the Structure of Reasons, *Philosophical Perspectives* 13: 57-89
- Conee, E. (2005) Contextualism Contested, in: M. Steup and E. Sosa, eds., *Contemporary Debates in Epistemology*, Malden (MA), Blackwell: 47-57
- DeRose, K. (1995) Solving the Skeptical Problem, *Philosophical Review* 104: 1-52
- Dretske, F. (1970) Epistemic Operators, *Journal of Philosophy* 67: 1007-23. Reprinted in his *Perception, Knowledge, and Belief: Selected Essays*, Cambridge, Cambridge University Press: 30-47

- Dretske, F. (1971) Conclusive Reasons, *Australasian Journal of Philosophy* 49 (1): 1-22. Reprinted in his *Perception, Knowledge, and Belief: Selected Essays*, Cambridge, Cambridge University Press: 3-29
- Dretske, F. (1981) The Pragmatic Dimension of Knowledge, *Philosophical Studies* 40: 363-78
- Dretske, F. (2005) The Case Against Closure, in: M. Steup and E. Sosa, eds., *Contemporary Debates in Epistemology*, Malden (MA), Blackwell: 11-26
- Feldman, R. (1999) Contextualism and Skepticism, *Philosophical Perspectives* 13: 91-114
- Firth, R. (1978) Are Epistemic Concepts Reducible to Ethical Concepts? In: A. Goldman and J. Kim, eds., *Values and Morals*, Dordrecht, D. Reidel: 215-29
- Goldman, A. (1976) Discrimination and Perceptual Knowledge, *Journal of Philosophy* 73 (20): 771-91
- Harman, G. (1986) *Change in View*, Cambridge (MA.), MIT Press
- Hawthorne, J. (2004) *Knowledge and Lotteries*, Oxford, Oxford University Press
- Lewis, D. (1996) Elusive Knowledge, *Australasian Journal of Philosophy* 74: 549-67. Reprinted in his *Papers in Metaphysics and Epistemology*, Cambridge, Cambridge University Press, 1999: 418-45
- MacFarlane, J. (2000) *What Does It Mean to Say that Logic is Formal?* PhD Dissertation, University of Pittsburgh
- Nozick, R. (1981) *Philosophical Explanations*, Cambridge (MA.), Harvard University Press
- Pritchard, D. (2010) Relevant Alternatives, Perceptual Knowledge and Discrimination, *Noûs* 44 (2): 245-68
- Pritchard, D. (2012) *Epistemological Disjunctivism*, Oxford, Oxford University Press
- Pritchard, D. (2016) *Epistemic Angst: Radical Scepticism and the Groundlessness of Our Believing*, Princeton, Princeton University Press
- Pryor, J. (2000) The Skeptic and the Dogmatist, *Noûs* 34: 517-49
- Pryor, J. (2001) Highlights of Recent Epistemology, *British Journal for the Philosophy of Science* 52 (1): 95-124
- Radford, C. (1966) Knowledge—By Examples, *Analysis* 27: 1-11
- Schaffer, J. (2005) Contrastive Knowledge, in: T. S. Gendler and J. Hawthorne, eds., *Oxford Studies in Epistemology* 1, Oxford, Oxford University Press: 235-71
- Schaffer, J. (2007) Closure, Contrast, and Answer, *Philosophical Studies* 133: 233-55
- Stine, G. C. (1976) Skepticism, Relevant Alternatives, and Deductive Closure, *Philosophical Studies* 29: 249-61
- Vogel, J. (1990) Are There Counterexamples to the Closure Principle? In: M. Roth and G. Ross, eds., *Doubting: Contemporary Perspectives on Skepticism*, Dordrecht, Kluwer: 13-27
- Vogel, J. (1999) The New Relevant Alternatives Theory, *Philosophical Perspectives* 13: 155-80
- Williams, M. (1991) *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*, Oxford, Blackwell
- Williamson, T. (2000) *Knowledge and Its Limits*, Oxford, Oxford University Press
- Wright, C. (2003) Some Reflections on the Acquisition of Warrant by Inference, in: S. Nuccetelli, ed., *New Essays on Semantic Externalism and Self-Knowledge*, Cambridge (MA.), MIT Press: 57-77