CLOSURE INTUITIONS AND RESTRICTION

Intuiciones y restricción del Principio de Cierre

Leandro De Brasi
Universidad Alberto Hurtado, Chile

Abstract

In this article I consider some alleged intuitive costs concerning the denial of the full generality of the Principle of Closure for knowledge. Usually philosophers dismiss such denial as highly counter-intuitive but I argue that, at least with regard to the alleged costs here considered, this is wrong: given our folk-intuitions, there are no such costs. So a fallibilist who seeks to halt the closure-based sceptical argument can restrict the principle with no such intuitive costs.

Keywords: Closure Principle, Transmission Principle, Fallibilism, Scepticism, Abominable Conjunction.

INTUICIONES Y RESTRICCIÓN
DEL PRINCIPIO DE CIERRE

Resumen

En este artículo, considero algunos supuestos costos intuitivos relativos a la negación de la generalidad del Principio de Cierre para el conocimiento. Usualmente los filósofos descartan tal negación como altamente contra-intuitiva pero argumento que, por lo menos en relación a los supuestos costos aquí considerados, esto es incorrecto: dadas nuestras intuiciones folk, no hay tales costos. Por lo tanto un falibilista que busca detener el argumento escéptico basado en el principio de cierre puede restringir el principio sin sufrir esos supuestos costos intuitivos.

Palabras Clave: Principio de Cierre, Principio de Transmisión, Falibilismo, Escepticismo, Conjunción Abominable

Leandro De Brasi es Profesor Auxiliar de la Universidad Alberto Hurtado, Santiago de Chile. Es PhD Philosophy, King's College London; MPhil Philosophy of Psychology, University of London; MSc Philosophy of Mental Disorder, King's College London; BA Philosophy, University of London. Entre sus áreas de interés se encuentra: espistemología, filosofía de la mente, filosofía de la psicología y psiquiatría.

Dirección electrónica: ldebrasi@uahurtado.cl

CLOSURE INTUITIONS AND RESTRICTION

Leandro De Brasi
Universidad Alberto Hurtado, Chile

In this article I consider some alleged folk-intuitive costs concerning the denial of the universality of the Principle of Closure for knowledge that are usually raised by philosophers to straightforwardly dismiss it. It is often claimed that any theory that entails its denial must be seriously flawed since it seems "utterly uncontentious" that knowledge is closed under competent deduction (Pritchard, 2013, p. 68). I suggest however that, at least with regard to the alleged pre-theoretical costs here considered, this is wrong. I argue that, given our folk-intuitions, there are no such costs. So a fallibilist who seeks to accommodate our strong presumption in favour of the possibility of knowledge and so halt the closure-based sceptical argument can restrict the principle with no such folk-intuitive costs.

To be clear, my aim isn't to show that knowledge is possible but to argue that, to stop the closure-based sceptical argument, a fallibilist can restrict the principle with no folk-intuitive costs with which this move is usually associated. So here I am only interested in the general and particular intuitions of the folk. Of course one might have excellent theoretical reasons for denying the principle's universality that one thinks outweigh the alleged pre-theoretical reasons here considered for accepting it. So one might not feel compelled to answer to these intuitive objections. However, philosophers do use those alleged pre-theoretical reasons to quickly discard the denial and I want to show that this is not a legitimate move since our folk-intuitions are in line with a restriction of the principle.

The article proceeds as follows. I first introduce a relevant alternative version of a fallibilist approach. Given this background, I consider some alleged costs that this approach is meant to have when combined with the denial of the principle's universality, so to avoid reintroducing far-fetched error-possibilities. I start with the odd-sounding conjunctions with which it is compatible. Exploiting some Gricean maxims, I explain pragmatically why these conjunctions seem odd to the folk without recourse to the Principle of Closure. Still some might suggest that it is the intuitiveness of this unrestricted principle that is responsible for this impression of, as some would say, abomination. So next I consider the folk-intuitiveness of the Principle of Closure, which seems to derive from a general intuition that supports more directly a related but stronger Principle of Transmission. Once I introduce the major difference between the principles, I examine their relation and how this general folk-intuition supports the Principle of Closure without supporting its universality. Then I consider particular folk-intuitions that suggest restrictions to both principles and offer, given the fallibilist framework provided, a plausible epistemological explanation, based on the discriminatory nature of knowledge-yielding procedures, of how these principles can fail in the way our intuitions suggest. Finally I conclude that this fallibilist approach that restricts the Principle of Closure can both accommodate and explain our folk-intuitions about it.

Fallibilism and Relevant Alternatives

There is a strong commonsensical presumption in favour of the possibility of knowledge. Non-sceptical accounts of knowledge normally accept, in order to accommodate this possibility, fallibilism. This is the thesis that not all possibilities of error need to be eliminated in order to know that *p* (Cohen, 1988; Lewis, 1996). We don't need to eliminate *all* error-possibilities: to meet all conceivable epistemic challenges to *p*. In other words, we don't need to rule out any imaginable way for the given proposition to be false. A common idea instead is that we only need to meet those challenges that are *relevant*.[1] The idea is that only relevant alternatives, i.e. significant possibilities of error, need to be eliminated. Another common idea is that significant error-possibilities are possibilities that aren't far-fetched and unlikely *given our worldview*. Given this, when determining,

[1] Pritchard (2013, p. 65) talks about the *core relevant alternates intuition*: "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 […] *relevant*." He also suggests "the natural motivation for fallibilism [...] is the intuition that in order to know one only needs to be able to rule out the error-possibilities that are […] *relevant*" (2005, p. 35).

say, whether something is a goldfinch, we don't need to show that it isn't a stuffed goldfinch, unless our worldview renders that error-possibility a significant one. This then is in line with our folk-intuitions, say: we don't think that the detective needs to rule out the possibility that Queen Elizabeth II stole my car.

So, according to this particular fallibilist approach, the knowledge-yielding procedures eliminate only the relevant alternatives that challenge the truth that p (and not all possible challenges).[2] And here I take it that whether a given challenge is relevant with regard to p depends on our worldview: more precisely, on whether we reasonably take such a challenge to represent a likely error-possibility (Rysiew, 2001, p. 488, 2005, p. 54). In other words, if some possibility is reasonably taken not to significantly change the truth-conduciveness of a knowledge-yielding procedure, such possibility is discounted as irrelevant. This is then the sense of 'relevant' that we are working with and it is this sense that explains why we don't think we need to eliminate the possibility of the bird being a stuffed goldfinch when considering whether it is a goldfinch. And importantly, this is also why we don't think we need to rule out sceptical possibilities, such as evil demons and brains-in-vats scenarios. So this view matches our pre-theoretical folk-intuitions, as we would prefer (Greco, 2000, p. 205).[3]

The subject then is to meet all the epistemic challenges that are reasonable with regard to p (given her worldview). And these are the challenges that are dealt with by exploiting the relevant knowledge-yielding procedures, or, as J.L. Austin says, the various "recognized procedures (more or less roughly recognized, of course), appropriate to the particular type of case" (1970, p. 87). It is only the epistemic demands that are *enough to prove p* that ought to be met (Austin, 1970, p. 85). But,

---

[2] Hendricks (2006) refers to this kind of approach as *forcing*, where error-possibilities are forced out for being too remote or speculative and so not being considered.

[3] Or so I assume here given that my opponent in this article exploits alleged pre-theoretical folk-intuitions to quickly discard the denial of the Principle of Closure. But let me make a brief related point. Of course there might be an issue as to whether those intuitions should really be taken into consideration. Experimental philosophers have for a while now been suggesting that intuitions, folk or otherwise, are *not* reliable data for theory construction (e.g. Alexander *et al.*, 2010; Weinberg *et al.*, 2001). But, regardless of whether experimental philosophers are onto something (and certainly it isn't clear that they are—e.g. Boyd and Nagel, 2014; Turri, 2013), we needn't here worry about this challenge since again the opponent takes it that, for some reason (maybe because they take it to be reliable data or a starting point that grounds the inquiry to commonsense, from which our inquiries cannot deviate too far at the expense of changing the subject), this data has a role to play in theory construction.

"Enough is enough: it doesn't mean everything. Enough means enough to show that [within reason] it 'can't' be anything else, there is no room for an alternative, competing, description of it. It does *not* mean, for example, enough to show that it isn't a *stuffed* goldfinch" (1970, p. 84).

Now, this fallibilism is welcome since it isn't clear we could have much knowledge if infallibilism were to be adopted: conclusive grounds, where all possible epistemic challenges to *p* are ruled out, aren't readily available for creatures like us. In most cases, our grounds seem to be logically consistent with the truth of not-*p*: that is, our grounds normally don't seem to entail or necessitate the truth of *p* (consider, for example, Barry Stroud's (1984) plane-spotters who identify the planes on the basis of inconclusive information, and so is our evidence, when identifying goldfinches and zebras). Indeed, infallibilism seems to rule out knowledge in many different domains, such as scientific and historical ones, that we would normally claim to have it. As David Lewis says, "uneliminated possibilities of error are everywhere. Those possibilities of error are far-fetched, of course, but possibilities all the same. They bite into even our most everyday knowledge" (1996, p. 549). And, since such an infallibilistic position would conflict with our folk-intuition about the possibility of knowledge in different domains, fallibilism seems preferable.[4] Indeed, its continuing widespread acceptance in epistemology can be seen as a result of this piece of commonsense (e.g. Cohen, 1988, p. 91; Vahid, 2008, p. 325).[5]

Having said that, if we accept a Principle of Closure [C] that says, roughly, that if one knows that *p*, competently deduces *q* from *p* and believes that *q*, then one knows that *q*,[6] [C] reintroduces the need to eliminate the remote (sceptical) possibilities in order to know ordinary propositions. After

---

[4] Of course if we are to answer Stroud's complaint that we "cannot conclude simply from our having carefully and conscientiously followed the standards and procedures of everyday life that we thereby know the things we ordinarily claim to know [...given that the] admitted fact that we do not insist on eliminating [the sceptical] possibility in everyday life does not show that we do not need to eliminate it in order to know" (1984, p. 69), we would need to independently motivate the above fallibilist approach. But as mentioned before, my aim here isn't this. I don't aim to show that knowledge is possible but show that, to stop the closure-based sceptical argument, this fallibilist can restrict closure with no folk-intuitive costs that this move is usually associated with.

[5] I here take the folk-intuitive data to be part of commonsense.

[6] Where *p* can stand for [*p* and (*p*→*q*)]. Baumann (2011, p. 599) refers to [C] as expressing "the core of the idea of knowledge closure currently discussed." Cf. Dretske (2010, p. 133): "Closure is the principle that one knows, or at least one is automatically positioned to know [...], all the known logical consequences of things one knows." But we won't be concerned with weaker and more vague being-positioned-to-know closure principles or stronger and false closure of knowledge under entailment principles.

all, given [C], if I don't know that (*q*) I'm not a brain in a vat (BIV), then I don't know that (*p*) I have a hand.[7] In other words, [C] reintroduces the problem we were meant to avoid by adopting fallibilism: having to eliminate extraordinary possibilities in order to know ordinary propositions.

Although this problem could be avoided if [C] were not to hold unrestricted, most philosophers don't seem keen on exploiting this strategy. Indeed, it is usually claimed that [C], or something very much like it, is highly intuitive or extremely plausible and the cost to deny it is too much for any theory to pay (e.g. Hawthorne, 2004; Pritchard, 2013; Williamson, 2000). This is partly why the closure-based sceptical argument can be and is often recast as a paradox, where we (putatively) have three individually intuitive claims that seem jointly inconsistent (Cohen, 1988): (a) I know that I have a hand (*p*), (b) I don't know that I am not a BIV (*q*), and (c) If I know that *p*, I know that *q*. Moreover, the denial of [C]'s full generality would seem to commit us to the possibility of "abominable conjunctions" (DeRose, 1995), such as that we can know that I have a hand and not know that I am not a BIV.

Below I consider these alleged costs concerning the denial of [C]'s full generality. I argue that they aren't costs and that the above fallibilist who seeks to accommodate the possibility of knowledge can deny [C]'s universality and indeed ought to do so if we are to capture our folk-intuitions (as the opponent seems to assume—see fn.3). So let us start by considering the aforementioned abominable conjunctions and later on the intuitiveness of [C].

Abominable Conjunctions

Fallibilism seems to commit us to abominable conjunctions since it seems that, say, when looking at the zebra, I can know that it is a zebra but not know that it is a painted mule (given that the ordinary procedure usually exploited to determine that it is a zebra doesn't rule out the possibility of it being a painted mule). So taking for granted that the conjunctions seem abominable, I want to examine where this impression of abomination comes from. Let us then evaluate under which circumstances, if any, we would ordinarily hold such a conjunction. Consider the following conversation between two adults in the zoo:

A: Do you know what animal that is?
B: Yes, it's a zebra, just look at the stripes.
A: How do you know that it isn't a painted mule?
B: Well, I don't.

---

[7] Schematically, the sceptic argues as follows: (1) Not-$K_{Not\text{-}BIV}$, (2) $K_{Hands} \rightarrow K_{Not\text{-}BIV}$ ∴ Not-K_{Hands}; where (2) derives from [C].

A: So, do you really know that it's a zebra?

I take it that B's answer to A's last question could be that he does know that it is a zebra if B is aware that A doesn't have any concrete reasons for introducing the painted-mule possibility (and their sharing a worldview helps determine this). In other words, even if the painted mule possibility is *salient* in the conversation, given their shared worldview and that no concrete reason in support of it is given or thought to be available in this particular case, such possibility isn't considered *relevant* (that is, the shared background information makes the painted-mule possibility insignificant). Salient possibilities needn't be regarded relevant.

But if that is the case, wouldn't B also say, contrary to our dialogue, that he does know that it isn't a painted mule? After all, B could infer that he also knows that it isn't a painted mule *given the background information*. That is, he can use that information (say, that zoos don't commit such deceptions, that the consequences would be severe, that such a deception isn't likely to last, etc.) to make an inference to the best explanation (IBE) or an inductive inference that the animal isn't a painted mule (although such background information isn't always available; consider again Stroud's (1984) airplane-spotters).[8] In this case then we would be applying some inferential procedure (as opposed to, say, a perceptual one: a procedure exploiting perceptual capacities) to know that it isn't a painted mule. So, given there is no reason to take the possibility seriously, we would normally think that we can extend our knowledge in that way (viz. via an IBE or inductive inference) in this case.[9] This then doesn't turn out to be a case where the abominable is held, since we would end up attributing knowledge that it is a zebra and that it isn't a painted mule.

Nevertheless, we *qua* folk would normally withdraw the knowledge claim and perhaps even reverse it, after being presented with a possibility that we can't eliminate. More precisely, we usually seem to withdraw but not reverse the claim, if not clear as to whether the introduction of the

---

[8] That is, we can in this case extend our knowledge by means of non-deductive reasoning. I'm assuming that IBE isn't a disguised inductive inference. In an IBE, one infers a hypothesis that explains a given set of data better than competing hypotheses (which is what, say, the detective normally does, Sherlock Holmes aside, when confronted with the clues). Normally, features such as simplicity, economy and fit with background information are put forward as the kind of virtues that allow us to choose among alternatives. The inductive inference might instead go like this: All zoos so far haven't engaged in such a deception (the fact that there are no reports that zoos engage in such deception supports this general statement), so in all likelihood this zoo isn't engaged in such a deception.

[9] Assuming a knowledge norm of assertion, this allows us to explain cases such as Kvanvig's (2008, p. 477), where we assert the animals aren't "fancy *robotics*."

possibility is legitimate (i.e. based on concrete reasons), and we usually seem to withdraw and reverse it, if clear about the legitimacy of such introduction. So, given that we naturally think that knowledge requires elimination of significant (i.e. relevant) error-possibilities, the reason for the withdrawal is that we aren't sure whether the painted-mule possibility is a significant one and the reason for the reversal is that we think such possibility is significant.

Importantly, assuming a Gricean approach (where our talk exchanges are governed by something like the Cooperative Principle: make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged),[10] when people introduce error-possibilities, we naturally think that there is some concrete reason for the introduction given the Maxim of Relation (i.e. to be relevant—otherwise, there would be no point in doing so) and so naturally reverse the claim. As Austin would say, introducing such an error-possibility doesn't (normally) mean merely that you are a fallible human being: "it means that you have some concrete reason to suppose that you may be mistaken in this case" (1970, p. 98). After all, if the possibility isn't significant, the Maxim of Quantity (i.e. to be as informative as required) recommends that one doesn't raise it.

Applying this to our case, the painted mule possibility is now taken to be not merely a salient possibility but a relevant one that isn't eliminated by the (e.g. perceptual) procedure exploited to determine whether the animal is a zebra in the ordinary case, hence we naturally end up with a reversal. Once the painted-mule possibility is introduced in the above circumstances where one takes it to be a significant possibility, it would be conversationally infelicitous to continue claiming that one knows that it is a zebra since it implicates that one can rule it out. Indeed, unless one thinks there are no concrete reasons for the introduction of the possibility and even if it actually is an insignificant one, one doesn't take oneself to know given that a putative significant error-possibility hasn't been eliminated. In these circumstances we wouldn't hold the conjunction either. So, given the above, we wouldn't normally hold the abominable conjunction, whether we are clear or not as to whether the possibility is significant or not. And so this seems to explain why we as folk think it is an abominable conjunction (cf. Dretske, 2005).

Closure and Transmission

But, isn't the problem really that such conjunction violates the Principle of Closure for knowledge [C]? This principle, one might think, is what

---

[10] This is Grice's "rough general principle which participants will be expected […] to observe" (1989, p. 26).

makes the conjunction abominable, since it (or something like it) is meant to be intuitive. Since one knows that it is a zebra and that the animal being a zebra entails that it isn't a painted mule, one seems to also know that the animal isn't a painted mule. So it seems wrong to claim that one knows that it is a zebra but doesn't know that it isn't a painted mule. The counter-intuitiveness of the abominable conjunction, it might be claimed, is the result of our holding an intuitive and unrestricted Principle of Closure. Indeed, even those who, for theoretical reasons, deny the principle admit to its intuitive plausibility (e.g. Nozick, 1981, p. 205-6). But, what is the content of this "closure intuition" and is it a folk-intuition that, as a piece of commonsense, would be preferable to accommodate (cf. Kvanvig, 2008, p. 476)? What is the folk-intuitive data with regard to the principle?

Certainly the content of some such closure intuition isn't some specific formulation of the principle, say, [C] above. And if [C] were the content of the intuition, it would need to be true (otherwise there would be no problem in rejecting it), and [C] isn't since, say, one might cease to know p by the time one comes to believe q. Indeed, refinements to [C] (e.g. that one retains the knowledge that p) are required if we are to have a Principle of Closure that is true (Hawthorne, 2004, 2005). But as Robert Nozick (1981, p. 205) says: "We would be ill-advised, however, to quibble over the details of [C]. Although these details are difficult to get straight, it will continue to appear that something like [C] is correct" (see also Hawthorne 2004, p. 36). Below I suggest why something like [C] seems correct, but now we need to notice that it is difficult to reach a version of [C] that is clearly true.[11] And even if there is some specific true formulation of the principle, it isn't immediately intuited. So it seems that, whatever the content of our closure intuition is, it isn't some specific true principle.

We *qua* folk seem to have instead an intuition whose content seems to be something like the idea that *we can gain knowledge by competent deduction* (or "*that deduction is a good epistemic method*, or *that knowledge can be extended by deduction*;" Lawlor, 2005, p. 31). And John Hawthorne says

234

---

[11] Particularly if the principle is to give us a plausible closure-based sceptical argument. David and Warfield (2008, p. 159) argue there is a restriction issue, since the refinements force the sceptic to make assumptions (in order to reach their general conclusion), such as attributions of "lots of beliefs involving denials of lots of different particular sceptical hypotheses" to the folk, which are clearly implausible. And add that the more refinements are introduced to make [C] plausible, the more problematic beliefs are required (*ibid.*). Cf. Lawlor (2005), where the suggestion is that, since one has antecedent reason for believing that one doesn't know that one isn't a brain-in-a-vat and closure doesn't apply if there is evidence against p (in this case, the ordinary proposition), [C] together with the refinement suggested above can't generate the paradox. See also Silins, 2005, pp. 89-91.

that [C], together with some refinements (as the one pointed out above), is a "more satisfying development of the closure intuition. The core idea behind closure is that we can add to what we know by performing deductions on what we already know" (2005, p. 29; see also Williamson, 2000, p. 117). Indeed, Jonathan Kvanvig identifies "the intuitive idea" behind [C] in the current debate as being "that knowledge can be extended by deduction" (2008, pp. 456, 474).

But this core general intuition seems to support more directly a slightly different principle: a Principle of Transmission for knowledge [T]. This is a related but stronger principle that states that knowledge is transmitted across competent deduction: if one knows that *p*, competently deduces *q* from *p* and believes that *q*, then *one comes to know that q on the basis of the deduction*.[12] The main difference between (some version of) [C] and [T] is that, while [T] clearly states that one knows that *q on the basis of the deduction*, [C] is silent as to where the knowledge that *q* comes from and, in fact, it needn't be arrived at by the deduction (Baumann, 2011, p. 600; Brown, 2004, p. 242).[13] [T] is stronger than [C] because [T] is more specific than [C]: [T] says everything [C] says and it adds that the conclusion is known in virtue of one knowing the premises and competently deducing it. So, given that in cases where a deduction fails to transmit knowledge, one can still know the conclusion (*q*) through some *other* means, counter-examples to [T] aren't counter-examples to [C]. But importantly for present purposes, if [T] is true, [C] is too since deduction is a means of coming to know that *q*.

So the "core idea" that we can gain knowledge by deduction makes [T] intuitive, which in turn makes [C] intuitive. This general folk-intuition then indirectly supports [C] via [T] (which allows knowledge to be extended through deduction), hence the intuitiveness of *any* such [C], as Nozick points out. And this intuition can also explain why we think that we should be "very reluctant" to reject [C] since: "If we reject it, in what circumstances can we gain knowledge by deduction?" (Williamson, 2000, p. 118; we answer this question below). After all, if [C] fails, [T] does too since deduction represents

---

[12] Just as with [C], there are slightly different formulations of [T] in the literature (see e.g. Pritchard, 2013, p. 75), but this one will do for present purposes since it helps us to highlight the most significant difference between the principles, which allows us to understand their relation. But notice that I am setting aside the issue as to whether the principle requires *de novo* knowledge of the conclusion: i.e. it requires the acquisition of knowledge *not previously had by any other means*. As we shall soon see, what matters to us is that the principle specifies *one means* by which that knowledge is present, regardless of whether this is the *only means* by which knowledge is present. For more on the distinction between the principles, see e.g. Davies (1998) and Wright (2000).

[13] I hereby ignore the parenthetical qualification.

one way of coming to know the conclusion that q that we can't come to know that q given [C]'s failure. So [C]'s failure threatens [T], which is in turn supported by the "core idea."

But we should notice that this intuition doesn't address the full generality of [T]: it merely says that we *can* gain knowledge by deduction. The intuition is in fact compatible with restrictions to [T]. As Kvanvig says, the claim that "deduction is always and everywhere knowledge-extending does not follow from the obviously correct claim that deduction is a way of extending knowledge" (2008, p. 474). So it is consistent with commonsense for both [T] and [C] to be restricted but not altogether rejected, unless [T]'s restrictions are such that the possibility of gaining knowledge by deduction is more or less limited than what we normally take it to be (given consideration to particular cases). But if this is so, the denial of the full generality of [C], which borrows its intuitiveness from the "core idea," needn't count as not capturing the commonsense data.

Of course, none of this shows that [C] should be restricted. But if the above is correct, it puts pressure on the claim that the "abominable conjunctions" are counter-intuitive *for the folk* because of [C], since at least the general folk-intuition doesn't commit us to [C]'s universality. And as we shall next see, particular folk-intuitions concerning [T] support a restriction of [T] in extraordinary cases (such as the zebra case), hence those intuitions failing to support [C] in those cases too. And although we shouldn't forget that counter-examples to [T] aren't counter-examples to [C] (so [C] can still retain its full generality even if our intuitions suggest restrictions to [T]), the intuitive case for the above explanation of the abomination of the conjunctions is surely weakened. So let us introduce next the case for an intuitive restriction to [T] and consider afterwards whether an analogous one applies to [C] too.

Easy Knowledge and Discrimination

When confronted with some cases of deductive reasoning instantiating [T]'s antecedent, we as folk have the intuition that we can't claim to gain knowledge of the conclusion (Cohen, 2005). In these cases, [T] seems wrong since it allows for knowledge of certain propositions via the deduction *far too easily*. For example, when determining (in ordinary cases) whether the animal is a zebra via a perceptual procedure mainly involving looking at it, we don't consider whether it is a painted mule, so the claim that we know via a deduction that it isn't a painted mule (based on our knowledge that it is a zebra and that if it is a zebra then it isn't a painted mule) seems in this case inappropriate. Indeed, if one says, when answering the legitimate question "How do you know it's not a painted mule?", "I know it's a zebra

[via ordinary means]; so, deductively, I know it's not a painted mule," then, as Kvanvig says, a "sense of shame would be appropriate at putting on such airs" (2008, p. 477). So some restriction to [T] doesn't only seem compatible with the commonsense data, but also required by it.

One might at this point suggest that the problem with the reasoning isn't [T] but the premise claiming knowledge that it is a zebra. That is, the reasoning seems suspect because we don't actually know that it is a zebra. But this, I take it, isn't what the folk would say. That is, we normally claim to know so by looking at the animal and seeing its stripes. One might anyway suggest that the problem is that it seems as if one doesn't know that it is a zebra: after all, the possibility of it being a painted mule is salient. But, as seen, when the possibility isn't taken to be significant, we won't normally either withdraw or reverse the knowledge claim. In these cases, where the possibility of being a painted mule is salient but not taken to be relevant and so not thought as having to be eliminated in order to know that (p) it is a zebra, we still think that there is something amiss in the above reasoning when concluding that one knows that (q) it isn't a painted mule by means of a deductive inference based on the fact that p entails q, and precisely because the possibility of being a painted mule hasn't been eliminated in order to know that q. And although in those cases, we *can* know through an IBE or induction (given our background information) that the animal isn't a painted mule, the issue here concerns the *gain of knowledge through deduction*. Hence, there seems to be a problem with [T]. And if we are right about it, this is folk-intuitive data that we would prefer to accommodate (by the lights of the opponent).

Indeed, it seems clear that such reasoning won't work, since if we were looking for a painted mule that looked like a zebra, we would look for an animal that looked like a zebra (i.e. the perceptual procedure used to know that it is a zebra in the example above) and then further investigate whether it is a zebra or a painted mule (say, by checking for paint). So it isn't clear in this case how we can deduce from our knowledge that it is a zebra via the ordinary means of looking at the animal that we *know* that it isn't a painted mule (Cohen, 2005, p. 424; Wright, 2003, p. 60). Again, in that case, if we were asked how we knew that it isn't a painted mule, we would find it absurd to reply that we do so *because* it looks like a zebra ("Look at the stripes!") and if it is a zebra then it isn't a painted mule, since it doesn't seem that we are in a position to know. It seems that we would know *too easily* that the painted mule hypothesis is false. This is of course a version of the "problem of easy knowledge" (Cohen, 2002, 2005), where some basic knowledge interacts with a deductive rule to allow us to know certain propositions far

too easily given the circumstances.[14] Here we are exploiting this problem to point out that [T] seems to require some restriction and so that the denial of the full generality of [T] actually counts as capturing the intuitive data.

Therefore the above general folk-intuition is consistent with a restriction on the class of known propositions that deduction can be used on to gain knowledge and, I suggest, a restriction seems required when considering our particular folk-intuitions concerning specific cases, as the above exemplifies. Now, according to this data, [T] doesn't apply to all known propositions, but let me here suggest an explanation as to why this is so.

The knowledge-yielding procedures are truth-discriminating: that is, they discriminate truths from falsehoods. So the procedures provide us with the capacity to distinguish competing state of affairs. This is easily seen in the zebra case: when exploiting the perceptual procedure, one distinguishes there being a zebra from, say, there being a giraffe. And this ability to tell the difference between competing states of affairs is an essential element of what makes a procedure truth-conducive. So, as one would expect, this notion of discrimination is central to reliability (Goldman, 1976, 1986). As Alvin Goldman says, "To be reliable, a cognitive mechanism must enable a person to discriminate or differentiate between incompatible states of affairs" (1976, p. 771). He motivates this claim by means of the famous fake-barn case, where one can't distinguish the real from the fake barns and so one doesn't seem to know. That is, he exploits our intuition that knowing that p requires the ability to distinguish p from state of affairs where p is false.

This ability to discriminate is central to our understanding of the procedures as being truth-conducive and it is in line with the idea that knowledge requires a discriminative capacity: the capacity to discriminate truth from falsehood. If it was lacking, one would be achieving the truth accidentally and not have knowledge (McGinn, 1999). Consequently, "A knowledge attribution imputes to someone the discrimination of a given state of affairs from possible alternatives" (Goldman, 1976, p. 772). Indeed, this ability to discriminate is what provides the proper connection to the world that knowledge requires and we desire. So, this requirement is motivated by the fact that "knowledge is a matter of responsiveness to the way the world is" (Roush, 2005, p. 122) and discrimination is the natural option that allows

---

[14] Basic knowledge is knowledge that one has prior to knowing that its source is legitimate. We won't deal with the other problem of easy knowledge: "bootstrapping," where easy knowledge is obtained via track-record arguments that exploit deductions (Vogel, 2000; Cohen, 2002; Van Cleve, 2003).

us to achieve such responsiveness.[15] Knowledge requires a proper connection to the fact that makes the target proposition true, where a proper connection is a connection that allows us to discriminate such proposition from other state of affairs where it is false. And I suggest it is because of this constraint that the deductive rule fails, in certain cases, to transmit. Let me illustrate.

I come to know, say, that I have hands via a perceptual procedure that doesn't eliminate the BIV alternative since it is irrelevant (given our worldview). So when I infer, on the basis of having hands entailing not being a BIV, that I am not a BIV, this belief isn't knowledge, even if true, because in these cases even if the BIV possibility isn't relevant for the having-hands claim, it certainly is the kind of state of affairs that needs to be discriminated for the not-being-a-BIV claim, given that knowledge involves a discriminative capacity. More specifically, there is no discrimination of the target proposition in this particular piece of reasoning. So there is no discriminative capacity in place with respect to the target proposition. When the procedure via which we come to know p doesn't discriminate q then we can't know via deduction that q, since we would fail to discriminate q from other state of affairs where not-q. That is, since in order to know the conclusion one has to discriminate a possibility that isn't relevant when coming to know the premises, one doesn't come to know the conclusion *via the deduction*. So if there is an expansion of discriminatory power between the premises and the conclusion, the deduction fails to transmit.

Now, ordinary procedures (perceptual or otherwise) don't discriminate extraordinary propositions since these aren't relevant (that is, the truth-discriminating power of the ordinary procedures is limited to certain significant possibilities given our worldview). So we can expect a [T] failure when some such ordinary procedure is employed to know the ordinary proposition that entails the extraordinary conclusion. For example, one can't know *by means of the deduction* that it isn't a painted mule from one's knowledge that it is a zebra by means of looking at it since it

---

[15] Sensitivity (which is here understood as the satisfaction of the following subjunctive: not-p→not-B(p), which in a crude but intuitive reading states that "if p were false, S wouldn't believe that p" and only close possible worlds are relevant to its evaluation) is another way of being responsive to the world. But, Roush argues, safety fails to be so (where safety is understood as the satisfaction of the following subjunctive: B(p)→p, which in a crude but intuitive reading states that "if S believed that p, then p would be true" and, again, only close possible worlds are relevant to its evaluation). It fails because "it gets the direction of fit wrong for what knowledge is" (2005, p. 121). Regardless of this, neither of them seems the natural option and they don't seem required to account for the non-accidentality of knowledge (De Brasi, *forthcoming*). Anyway, I obviate them here since my aim is merely to offer a plausible explanation of the restriction that commonsense seems to suggest.

doesn't take into account the extraordinary possibility. After all, as Fred Dretske says,

> "Our ways of discovering P are not necessarily ways of discovering what we know to be implied by P. From the fact that you know that P implies Q, it does not follow that you can see (smell, feel, etc.) that Q just because you can see (smell, feel, etc.) that P" (2005, p. 14).

So one can't come to know, say, that one isn't a BIV via deduction from one's knowledge that one has hands via some ordinary (e.g. perceptual) procedure. But there might be other ways in which one can come to know that one isn't a BIV anyway, such as an IBE (but notice, given the way the sceptical hypotheses are designed, not many procedures will be able to deal with them).[16] And so importantly if we were to achieve knowledge of an ordinary proposition *p* via an IBE that discriminates the sceptical situation *q* in order to determine that *p*, one can come to know via deduction that *q*. Having said that, the Moorean inference, from my knowledge that I have a hand via some ordinary procedure, is deeply counter-intuitive and our explanation allows us to embrace this folk-intuition that one doesn't come to know that one isn't a BIV by means of this inference.[17]

Similarly, there are other ways in which one can know that it isn't a painted mule, say by checking for paint or through an IBE as suggested above. Importantly, there are ways in which we can come to know that it isn't a painted mule via deduction if we know that it is a zebra through one of these (non-ordinary) procedures (given usual circumstances). After all, those are the kind of procedures that help us determine that it is a zebra by ruling out the possibility (among others) that it is a painted mule. Anyway, even if we can gain knowledge that it isn't a painted mule via deduction when employing certain discriminatorily powerful procedures to know that it is a

---

[16] The IBE contrasts the target proposition with competing hypotheses, and so discriminates it from other states of affair where it is false. Certainly philosophers have often suggested and sometimes attempted (though schematically) such IBEs and the hypothesis of not being a BIV seems to be a better explanation of the data than the alternatives. But, in order to make the inference, the folk wouldn't only need to be aware of the relevant data but also of the different competing hypotheses so to weigh their merits. Note that, when explaining the data, the negative hypothesis "not being a BIV" needs to be backed up by a specific ("real-world") hypothesis (see Bonjour, 1985; Vogel, 2008). So, although some *might* be able to know via IBE that one isn't a BIV, most people aren't. That is, given the requirements on performing such IBEs, their successful employment is significantly restricted. So we can expect the IBE to normally fail to provide the folk with the knowledge that one isn't a BIV (even if it is a legitimate means to know so). This can explain why the folk normally think that they don't know that they aren't BIVs: there is no procedure they can normally implement.

[17] Schematically, the Moorean argues as follows: K_{Hands}, $K_{Hands} \rightarrow K_{Not\text{-}BIV}$ $\therefore$ $K_{Not\text{-}BIV}$.

zebra, these procedures aren't the ones we normally need to use or usually use (given our worldview) in order to know that it is a zebra. And it won't be known through a deduction that exploits one's knowledge that it is a zebra when achieved through an ordinary procedure (since this will be an item of knowledge that isn't powerful enough to allow for the discrimination of the target proposition). So, although one can normally gain knowledge *via deduction* of ordinary propositions (say, that the animal isn't a giraffe) and even sometimes of extraordinary propositions (say, that the animal isn't a painted mule—as long as there is no expansion of discriminatory power), one cannot *always* do so.[18]

Anyhow, the cases where we don't achieve knowledge that *p* via some such discriminatorily powerful procedure are also the cases where, as seen above, it seems that we gain knowledge too easily: that we can get knowledge on the cheap. After all, as Colin McGinn says, "it can be easier to know *p* than *q* though *p* implies *q* (and not vice versa) because *q* requires more in the way of discrimination than *p*" (1999, p. 27; see also Goldman, 1986, p. 56). In these cases, [T] doesn't satisfy the above discriminatory constraint, since not-*q* isn't taken into account when considering *p*, hence not allowing us to discriminate *q* from state of affairs where it is false. Indeed, one can't come to know the conclusion when one's ordinary (loosely-described) evidence doesn't address it.

With [T] we are not merely transmitting truth, and since in these cases we aren't in a position to know the conclusion *via the deductive reasoning*, [T] fails. In these cases, it seems we would be helping ourselves to some unearned positive epistemic status required for knowing. And the lack of discriminatory power explains why we can't come to know the conclusion in some cases when engaging in the deductive reasoning. Now, this explanation isn't *ad hoc* given the independently motivated reason for this discriminatory constraint. Indeed, as mentioned, it is the natural way for us to think of procedures that are responsive to the world. And so an unrestricted [T] fails in some cases because one can have the positive epistemic status required for knowledge in the case of the premise but not the conclusion, based on the deductive reasoning, due to a discriminatory deficit. But we can see that if there is no discriminatory deficit, [T] holds. And this is exactly so in the uncontroversial cases where we can extend our knowledge *through deduction* (say, where we move from knowing that it is a zebra to the deduced claim that it isn't a giraffe).

---

[18] I have been ignoring the further and trivial possibility of [T] failure in cases in which one already knows the conclusion *if* one is to understand the principle as requiring *de novo* knowledge of the conclusion (see fn.12). So in cases where knowledge is present *before* the reasoning occurred, this particular [T] also fails.

The above then provides us with an explanation as to why we don't know in cases where the knowledge seems to be acquired too easily and do know in cases where it seems appropriate to extend our knowledge via deduction. So we can answer Ram Neta's question as to why it is that, in the relevant range of cases, we can't gain knowledge so easily (2005, p. 189): namely, there is a discriminatory deficit. So I can know that it is a zebra via ordinary means but not know that it isn't a painted mule via deduction from that item of knowledge. And I can know that I have hands via ordinary means but not know that I am not a BIV via deduction from that item of knowledge. But, although this stops us from having "easy knowledge," we saw we can still have knowledge via deduction in both cases. And more generally, we saw we can expand our knowledge deductively. As long as there is no lack of discriminatory power, we can gain knowledge through deduction. Again, from my knowing that it is a zebra and that if it is a zebra, it isn't a hippopotamus, I know that it isn't a hippopotamus (and so with goldfinches and cravens, and red and blue tables).

So I suggest that, consistent with commonsense (particularly, the general "core idea" and the particular folk-intuitions concerning "easy knowledge" cases), [T] is to be restricted and, given that [C]'s intuitive support rests on [T]'s, the folk-intuitive data doesn't support the full generality of [C] either. So the apparent abomination of the above conjunctions doesn't seem to rest, *for the folk*, on the intuitive appeal of [C] since it is silent with respect to [C]'s full generality. But this needn't mean that [C], unlike [T], doesn't enjoy such generality since counter-examples to [T] aren't counter-examples to [C] even if, as Dretske would say, appreciating the failure of [T] makes the failure of [C] "easier to swallow" (2005, p. 15). Next then I want to consider the universality of [C], since an unrestricted [C] would seem to reintroduce problems of easy knowledge as well as sceptical ones.

Closure Again

Recall that [C] doesn't specify the way we come to know the conclusion, so [T] being restricted doesn't show that [C] is too. Indeed, [C] doesn't concern any specific procedure by which we acquire the knowledge of the conclusion, it simply states that some attributions of knowledge about a subject are incompatible: more precisely, if the antecedent of [C] is satisfied (knowledge that *p*, etc.), you *must* attribute knowledge that *q* to the subject. So [C] merely states that we know the conclusion by whichever means. Indeed, we might know the conclusion by some other inferential method, such as an induction or IBE. And as pointed out above, we can know via some such procedure that, say, the animal isn't a painted mule, given our

background information. So even if the deduction fails to transmit knowledge in this case, we might still be able to know the conclusion. We can have [T] failure without [C] failure. But the fact that we *can* know via other means the conclusion (given the available background information) doesn't mean that we *always* know. The above concession doesn't help us establish [C]'s universality.

At this point, however, one might want to suggest that, in becoming aware of the entailment, one also becomes aware of an error-possibility for which one must "form a view about what entitles you to dismiss this possibility," since "becoming aware of an error-possibility that you know is incompatible with what you believe and being unable to rationally dismiss it is […] knowledge-defeating" (Pritchard, 2010, pp. 255-6, 261; see also Pritchard, 2013). And the background information is meant to provide the evidence to defeat this defeater. Now, if this is correct, we can easily appreciate that [C] holds in this case. After all, the evidence required to rationally dismiss the defeater to *p* eliminates the illicit expansion of discriminatory capacity otherwise normally present in such cases, since that evidence is meant to rule out the possibility that not-*q*.

Now, although this move doesn't commit one to the implausible view that one needs to dismiss these error-possibilities at all times in order to know ordinary propositions (2010, pp. 261, 265fn.19; 2013, p. 104fn.17), the move still seems to be too demanding. After all, as noticed, salient possibilities needn't be taken as relevant ones (i.e. as introducing a significant error-possibility; cf. Pritchard, 2010, pp. 260-1). The disagreement lies in the role salient possibilities might have as defeaters (Pritchard, 2013, pp. 79-80). The suggestion is that "if I hadn't considered the possibility before, I should now" and that even if I don't have any concrete reasons for thinking that this is a significant error-possibility, still "I should be able to rule it out" (2010, p. 262). But this seems too demanding since after all one doesn't have reasons to believe that the alternative represents a significant error-possibility (and neither it is one); so why would one consider eliminating this possibility? And why think that a possibility when first considered and classified as irrelevant provides us with the means to eliminate said possibility? After all, it seems that at least in some cases some people when first faced with some extraordinary possibility (say, the BIV hypothesis) will classify it as "off-the-wall" without that allowing them to rule it out.

Importantly, the above approach doesn't require the elimination of such a merely salient extraordinary possibility either. Indeed, ordinary procedures don't require us to eliminate extraordinary possibilities and only if we have reasons to regard an alternative as introducing a significant error-possibility,

there would be a need, given that the procedures' goal is to promote the truth, to eliminate it. In that case, such an alternative would count as a defeater. But if an extraordinary possibility is merely salient, we don't need to rule it out. Moreover, there is no need for one to be able to classify the salient possibility as extraordinary since, without reasons to back them up, no possibilities need to be taken into account, since ordinary possibilities are already taken into account in the procedure and extraordinary ones can be ignored. So, although the far-fetched possibility could certainly be ruled out when coming to classify it as such, this needn't and it is unlikely to be always the case, hence not eliminating the possibility of an illicit expansion in discriminatory power. That is, [C] is likely to fail in some cases for some people, and in particular for the folk when dealing with sceptical scenarios.[19]

So [C] doesn't hold in full generality. And significantly, by restricting [C] to those cases where there isn't a discriminatory deficit, we can avoid some unpalatable moves given the commonsense data. First, we can avoid the sceptical move: if we don't know that some extraordinary (sceptical or odd) hypothesis doesn't hold, we fail, given [C]'s universality, to know many ordinary propositions incompatible with it. But this seems wrong since we *qua* folk don't think we need to eliminate such hypotheses to have ordinary knowledge. Second, we can avoid the Moorean move: given that we know those ordinary propositions and [C]'s universality, we also know that those hypotheses don't hold. But this again seems wrong since we *qua* folk don't think we can normally eliminate those hypotheses. Third, we can avoid the rejection move: that is, the wholesale rejection of [C]. And as the "core idea"

---

[19] It seems that lottery scenarios provide us with more intuitive counter-examples to [C] (more precisely, with more folk-intuitions suggesting its restriction). After all, for a great many ordinary propositions (e.g. S will never be rich) that we intuitively think we know, there is some lottery proposition (e.g. S's ticket is a loser) that we intuitively think we don't know, although highly likely and even though "in each case the ordinary proposition entails the lottery proposition" (Hawthorne, 2004, p. 5). That is, although the lottery proposition is a logical consequence of the ordinary proposition that we seem to know, we don't seem to know the former; hence [C] seems to fail in these cases. But given our restriction to [C] and the nature of such restriction, this needn't concern us. To see this, consider the ordinary proposition that my car is in the parking lot. Assuming that I clearly remember where I parked it, we think that I know that my car is in the parking lot via a memory procedure. But, although this ordinary proposition entails the lottery proposition that my car hasn't been stolen and driven off (assuming there is only a small chance that this is so), we don't want to say that I know this lottery proposition. And this, according to the approach suggested, is to be expected since the procedure employed for knowing the ordinary proposition doesn't possess enough discriminatory power to allow us to know the conclusion. Moreover, since the possibility of the car being stolen is, *ex-hypothesi*, small, such possibility isn't relevant when coming to know the ordinary proposition. Anyway, [C] seems to fail since there doesn't seem to always be another way of making up for such discriminatory deficit.

behind [C] makes clear, this seems wrong. Although commonsense doesn't support [C]'s universality, hence allowing us to restrict [C], it is clear our general and particular folk-intuitions support a restricted [T], which suggests that [C] at least sometimes holds.[20]

Concluding Remarks

Just like [T], [C] seems to hold in some cases and fail in others. Still this can do justice to our "core idea" that knowledge can be gained by deduction. So this licenses deductive knowledge but importantly doesn't force us to attribute knowledge in easy knowledge cases (including those involving sceptical scenarios) since, in tune with our folk-intuitions, we can fail to know the conclusion due to a discriminatory deficit, as I have here suggested. So, as seen, we can explain, as it is sometimes demanded, "why it is that, for any particular piece of [ordinary] knowledge, it seems that we can inferentially expand it in some ways but not in others—even when the inference is the same across these cases" (Neta, 2005, p. 196).

Nonetheless, even if it can be true that one knows that it is a zebra but not know that it isn't a painted mule, we wouldn't, as seen, often hold or utter the above "abominable conjunctions" since we would normally think they are wrong and conversationally infelicitous. But even if we are unlikely to think or assert them, such propositions needn't be false.[21] Importantly, these conjunctions don't seem abominable due to their falsehood given an unrestricted [C]. Therefore, given that the above restriction move doesn't do violence to commonsense and we can explain why we consider the above conjunctions as abominable, we can conclude that a fallibilist approach of the sort here presented that restricts closure in order to accommodate the possibility of knowledge doesn't suffer from those oft-mentioned closure-related intuitive disadvantages.

All else being equal, there are no reasons for preferring non-closure-restricted accounts given the relevant folk-intuitions. Moreover, given the nature of the denial of [C]'s universality, we can halt the closure-based sceptical attack. More precisely, neither the anti-sceptic nor the sceptic can conclude by means of a closure-based argument that the other is wrong by

[20] A fourth move that we can also avoid is the "shifty" one, where the impression of inconsistency between our ordinary knowledge-claims, extraordinary knowledge-claims and [C] is explained away, roughly, in terms of the sensitivity of epistemic standards to non-traditional factors, such as the awareness of error-possibilities, practical stakes and the like (e.g. Cohen, 2000; Hawthorne, 2004). Although I have not touched upon these issues, notice a shifty approach isn't adopted here in order to capture the above commonsensical claims. Non-traditional factors play no role in determining the relevance of alternatives.

[21] Dretske (2005, p. 19) suggests this is in fact a general phenomenon.

rightly pointing out that we do have knowledge of some ordinary proposition (via ordinary means) or that we do lack knowledge of some extraordinary proposition, respectively.[22] The response to the sceptical paradoxes seems to be [C]'s restriction, which, to repeat and contrary to what is usually claimed, is in line with our folk-intuitions.

Referencias Bibliográficas

ALEXANDER, J., MALLON, R. & WEINBERG, J. (2010): "Accentuate the Negative", *Review of Philosophy and Psychology*, 1, pp. 297-314.

AUSTIN, J. L. (1970). *Philosophical Papers*, Oxford: Clarendon Press.

BAUMANN, P. (2011): "Epistemic Closure", en S. Bernecker & D. Pritchard (eds), *The Routledge Companion to Epistemology*, London: Routledge, pp. 597-608.

BONJOUR, L. (1985). *The Structure of Empirical Knowledge*, MA: Harvard University Press.

BOYD, K. & NAGEL, J. (2014): "The Reliability of Epistemic Intuitions", en E. Machery & E. O'Neill (eds), *Current Controversies in Experimental Philosophy*, London: Routledge.

BROWN, J. (2004). *Anti-Individualism and Knowledge*, MA: MIT Press.

COHEN, S. (1988): "How to be a Fallibilist", *Philosophical Perspectives*, 2, pp. 91-123.

COHEN, S. (2000): "Contextualism and Skepticism", *Philosophical Issues*, 10, pp. 94-107.

COHEN, S. (2002): "Basic Knowledge and the Problem of Easy Knowledge", *Philosophy and Phenomenological Research*, 65, pp. 309-29.

COHEN, S. (2005): "Why Basic Knowledge is Easy Knowledge", *Philosophy and Phenomenological Research*, 70, pp. 417-30.

DAVID, M. & WARFIELD, T. (2008): "Knowledge-Closure and Skepticism", en Q. Smith (ed), *Epistemology: New Essays*, Oxford: OUP, pp. 137-88.

DAVIES, M. (1998): "Externalism, Architecturalism, and Epistemic Warrant," en C. MacDonald, B. Smith & C. J. G. Wright (eds), *Knowing Our Own Minds: Essays on Self-Knowledge*, Oxford: OUP, pp. 321-361.

DE BRASI, L. (*forthcoming*): "Accidentality and Knowledge after Gettier", *Filozofija*.

DE BRASI, L. (*ms*): "Testimony and Value in the Theory of Knowledge."

DEROSE, K. (1995): "Solving the Skeptical Problem", *Philosophical Review*, 104, pp. 17-52.

DRETSKE, F. (2005): "The Case against Closure", en M. Steup & E. Sosa (eds), *Contemporary Debates in Epistemology*, Oxford: Blackwell, pp. 13-25.

---

[22] Moreover, if we were to successfully answer to Stroud's complaint (fn.4—see De Brasi (*ms*) to gain an idea as to how this can be done) and given that "an outstanding challenge for epistemology [is] to show how knowledge can be possible at all without being easy" (Van Cleve, 2003, p. 57), this fallibilist approach would both accommodate and explain the possibility of knowledge as well as the folk-intuitive data concerning easy knowledge. This would give us a very good reason for favouring it.

Dretske, F. (2010): "Fred Dretske", en J. Dancy, E. Sosa & M. Steup (eds), *A Companion to Epistemology*, 2nd edition, Oxford: Blackwell, pp. 130-4.

Goldman, A. I. (1976): "Discrimination and Perceptual Knowledge", *Journal of Philosophy*, 73, pp. 771-91.

Goldman, A. I. (1986). *Epistemology and Cognition*, MA: Harvard University Press.

Greco, J. (2000). *Putting Skeptics in Their Place*, Cambridge: CUP.

Grice, P. (1989). *Studies in the Way of Words*, MA: Harvard University Press.

Hawthorne, J. (2004). *Knowledge and Lotteries*, Oxford: OUP.

Hawthorne, J. (2005): "The Case for Closure", en M. Steup & E. Sosa (eds), *Contemporary Debates in Epistemology*, Oxford: Blackwell, pp. 26-42.

Hendricks, V. (2006). *Mainstream and Formal Epistemology*, Cambridge: CUP.

Kvanvig, J. (2008): "Closure and Alternative Possibilities", en J. Greco (ed), *The Oxford Handbook of Skepticism*, Oxford: OUP, pp. 456-83.

Lawlor, K. (2005): "Living without Closure", *Grazer Philosophische Studien*, 69, pp. 25-50.

Lewis, D. K. (1996): "Elusive Knowledge", *Australasian Journal of Philosophy*, 74, pp. 549-67.

McGinn, C. (1999). *Knowledge and Reality*, Oxford: Clarendon.

Neta, R. (2003): "Contextualism and the Problem of the External World", *Philosophy and Phenomenological Research*, 66, pp. 1-31.

Neta, R. (2005): "A Contextualist Solution to the Problem of Easy Knowledge", *Grazer Philosophische Studien*, 69, pp. 183-206.

Nozick, R. (1981). *Philosophical Explanations*, MA: Harvard University Press.

Pritchard, D. (2005). *Epistemic Luck*, Oxford: OUP.

Pritchard, D. (2010): "Relevant Alternatives, Perceptual Knowledge and Discrimination", *Nous*, 44, pp. 245-68.

Pritchard, D. (2013). *Epistemological Disjunctivism*, Oxford: OUP.

Roush, S. (2005). *Tracking Truth*, Oxford: Clarendon.

Rysiew, P. (2001): "The Context-Sensitivity of Knowledge Attributions", *Nous*, 35, pp. 477-514.

Rysiew, P. (2005): "Contesting Contextualism", *Grazer Philosophische Studien*, 69, pp. 51-70.

Silins, N. (2005): "Transmission Failure Failure", *Philosophical Studies*, 126, pp. 71-102.

Stroud, B. (1984). *The Significance of Philosophical Scepticism*, Oxford: Clarendon.

Turri, J. (2013): "A Conspicuous Art: Putting Gettier to the Test", *Philosopher's Imprint*, 13, pp. 1-16.

Vahid, H. (2008): "The Puzzle of Fallible Knowledge", *Metaphilosophy*, 39, pp. 325-44.

Van Cleve, J. (2003): "Is Knowledge Easy—Or Impossible? Externalism as the Only Alternative to Skepticism", en S. Luper (ed), *The Skeptics*, Aldershot: Ashgate Publishing Company, pp. 45-60.

Vogel, J. (2000): "Reliabilism Leveled", *Journal of Philosophy*, 97, pp. 602-23.

Vogel, J. (2008): "Internalist Responses to Skepticism", en J. Greco (ed), *The Oxford Handbook of Skepticism*, Oxford: OUP, pp. 436-50.

Weinberg, J., Nichols, S. & Stich, S. (2001): "Normativity and Epistemic Intuitions", *Philosophical Topics*, 29, pp. 429-60.

Williamson, T. (2000). *Knowledge and its Limits*, Oxford: OUP.

Wright, C. (2000): "Cogency and Question-Begging: Some reflections of McKinsey's Paradox and Putnam's Proof", *Philosophical Issues*, 10, pp. 140-163.

Wright, C. (2003): "Some Reflections on the Acquisition of Warrant by Inference", en S. Nuccetelli (ed), *New Essays on Semantic Externalism and Self-Knowledge*, MA: MIT Press, pp. 57-78.