Tim Kraft

Scepticism, Infallibilism, Fallibilism

Abstract

The relation of scepticism to infallibilism and fallibilism is a contested issue. In this paper I argue that Cartesian sceptical arguments, i.e. sceptical arguments resting on sceptical scenarios, are neither tied to infallibilism nor collapse into fallibilism. I interpret the distinction between scepticism and fallibilism as a scope distinction. According to fallibilism, each belief could be false, but according to scepticism all beliefs could be false at the same time. However, to put this distinction to work sceptical scenarios have to be understood as ignorance-possibilities, not as error-possibilities. To show that scepticism is not tied to infallibilism I reject the principle of unrestricted relevance according to which any error- or ignorance-possibility whatsoever is relevant. Instead I argue that the sceptic should distinguish between local and global ignorance-possibilities. Global ignorance-possibilities are relevant even though not all ignorance-possibilities are relevant. The result is a refined version of the Cartesian sceptical argument that avoids two traps other versions do not avoid.

1. *Introduction*

The aim of this paper is to clarify the relation between Cartesian[1] scepticism on the one hand and fallibilism and infallibilism on the other. I will present a sceptical argument that isn't tied to infallibilism and is substantially, not merely verbally, different from fallibilism. Thus, no matter whether fallibilism or infallibilism turns out to be true the sceptical problem remains. Note that my aim is neither to defend nor to repudiate either one of scepticism, fallibilism or infallibilism; my aim is solely to clarify what those positions amount to and to offer a sceptical argument that avoids two problems that some versions of it don't avoid.

According to scepticism, I don't know anything about the external world (or empirically). Arguments for scepticism are in danger of stumbling into either one of two traps. The sceptical argument claims that in order to know anything empirically one needs to know whether one is a brain in a

[1] Since I only discuss Cartesian scepticism in this paper, I will drop the adjective "Cartesian" from now on.

vat, but that nobody does know this. If one's conception of knowledge is infallibilist, it's (probably) true that one needs to, but can't know whether one is a brain in a vat. However, this invites the reply that the infallibilist assumption is too strong and demanding. The thesis that scepticism is tied to infallibilism and that this dependency is its Achilles' heel is among the most common replies to scepticism (Brueckner 2005, pp. 389 f.; 2011, pp. 85-87). If instead one's conception of knowledge is fallibilist, it's not only difficult to get the sceptical argument going, but it also looks as if the sceptic's claim is something the fallibilist can wholeheartedly agree with: the fallibilist (probably) agrees that nobody can conclusively rule out being a brain in vat, yet rejects the sceptical conclusion that therefore nobody knows anything empirically. Assuming fallibilism the disagreement between scepticism and anti-scepticism is bound to be merely verbal: both agree that we can't conclusively rule out being a brain in a vat and only disagree about whether we should be said to have knowledge. If the word "knowledge" was dropped, there would be no way to express the disagreement.[2] Thus the sceptic's dilemma is this: *either* her argument rests on too strong a conception of knowledge (i.e. infallibilism) *or* assuming fallibilism she ends up asserting something fallibilists are happy to concede (i.e. that no belief is safe from falsity).

A common strategy to deal with this problem is to investigate whether infallibilism really is too demanding. Infallibilism might turn out to be not so counterintuitive after all (Unger 1975; Klein 1981; Stroud 1984; Lewis 1996; cf. Williams 1991) or indispensable for solving other epistemological puzzles like the Gettier problem, the lottery paradox or concessive knowledge attributions (Howard-Snyder et al. 2003; Dutant 2007; Dodd 2011; Moon 2012; cf. Hetherington 2012). In this paper I will not enter this debate. Instead, I will concede that infallibilism is implausible for the sake of the discussion and argue that scepticism is neither tied to infallibilism nor only verbally distinct from fallibilism. To achieve this I argue for three theses. My first thesis is that the distinction between scepticism and fallibilism should be treated as a scope distinction. However, for that idea to work fallibilism and scepticism must be stated in terms of ignorance-possibilities, not in terms of error-possibilities; this is my second thesis. My third thesis is that the ensuing sceptical argument relies on epistemic principles that don't tie scepticism to infallibilism.

[2] The danger of ending up in a merely verbal dispute is vividly illustrated by Stroud: scepticism easily appears to be like the position that there are no doctors since nobody can heal any disease within two minutes (Stroud 1984, p. 40).

2. *Defining Scepticism, Infallibilism, Fallibilism*

In this section I provide an overview over the three isms involved. I begin with a short summary of scepticism and then move on to a longer discussion of definitions of fallibilism and infallibilism.

2.1. *Cartesian Scepticism and the Argument from Ignorance*

I will only deal with Cartesian scepticism in this paper, i.e. the variety of scepticism that relies essentially on sceptical hypotheses or scenarios like the brain in a vat scenario. There is some debate about how to reconstruct the scenario-based sceptical argument. Yet, what is uncontroversial is that the argument has to fit the following template, the *argument from ignorance* (cf. DeRose 1995, p. 1):

(1) I don't know whether I'm in *SH*.
(2) If I don't know whether I'm in *SH*, I don't know whether *P*.
(3) Therefore, I don't know whether *P*.[3]

In this template *SH* is a placeholder for the chosen sceptical hypothesis (evil demon, dreaming, brain in a vat, five minute world etc) and *P* is a placeholder for any of the target propositions (propositions about the external world, propositions about the past etc). The argument from ignorance is only a template because . . .

(4) . . . the sceptic has to spell out the details of *SH*,
(5) . . . the sceptic has to show why I *don't* know whether *SH*,
(6) . . . the sceptic has to show why I *need* to know whether *SH* in order to know any of the target propositions.

I will provide answers to these questions in the following sections. I won't deal with a load of additional questions, e.g. whether the argument from ignorance is knowledge-specific or is radical in the sense that it applies to other epistemic notions (justification, warrant, entitlement) as well, whether the sceptic needs to show that knowledge is impossible or whether she has already won when she can offer reasons for staying agnostic on whether knowledge is possible, whether the sceptic needs to rely on intuitive premises only etc. Furthermore, I only discuss the brain in a vat-scenario and sceptical arguments that target empirical knowledge (not knowledge about the past, other minds etc). Empirical knowledge is meant to cover all knowledge essentially depending on perceptual experiences (ex-

[3] I prefer "know whether" instead of "know that" to avoid factivity/presupposition problems potentially lurking around "know that" (cf. Wang and Tai 2010).

cluding experiential knowledge of one's own mental states). I prefer to understand scepticism as being targeted at empirical knowledge and not at knowledge of the external world because some beliefs about the external world aren't empirical beliefs (e.g. the belief that it's raining or it isn't) and, therefore, not threatened by scenario-based arguments.

2.2. *Defining Fallibilism and Infallibilism*

I now turn to fallibilism and infallibilism. The search for a suitable, not too idiosyncratic definition is complicated by at least three problems. The *first* problem is that it's already controversial what the two views are views about. Are they theses about belief, justification or knowledge? The *second* problem is that fallibilism/infallibilism are treated by some authors as claims about whether some condition is a necessary condition for belief, justification, or knowledge, whereas other authors treat them as theses about whether the respective necessary condition is ever met.[4] The *third* problem is that various conditions are supposed to be central for the distinction: a belief is infallible if it couldn't be false, but it's far from clear which kind of modality is at stake here.[5]

2.2.1. *Conditional and Unconditional (In-)Fallibilism*

The first problem is to specify what fallibilism/infallibilism are about. The basic use of "fallible" and "infallible" is to distinguish two kinds of belief, justifications or pieces of knowledge (note that the distinctions can be drawn without supposing that both classes are non-empty):

(B) *S*'s *belief* that *P* is infallible iff it's impossible that *S* believes that *P* and it's false that *P* (= C_B).

[4] Among recent surveys Reed (2012) focuses on the question what the necessary conditions of knowledge are, whereas Leite (2010) distinguishes two versions of fallibilism: the first is "the thesis that human beings are fallible about everything (or just about everything) they believe" (Leite 2010, p. 370), the second is a thesis about necessary conditions for justification or knowledge. Dougherty (2012) explicitly defines fallibilism as a thesis about the consistency of a conjunction:

> The epistemological doctrine of fallibilism, though, is about the consistency of holding that humans have knowledge while admitting certain limitations in human ways of knowing. (Dougherty 2012, p. 131)

In what follows I keep the two issues mentioned in the quotation apart, i.e. the issue of whether knowledge requires infallibility on the one hand and the issue of whether we ever are infallible on the other hand.

[5] A fourth problem, not dealt with in this paper, is whether belief in a *necessary* proposition can be fallible; cf. Reed (2012, p. 586) and the literature mentioned there.

(J) *S*'s *justification* for believing that *P* is infallible iff it's impossible that *S* is justified in believing that *P* and it's false that *P* (= C_J).

(K) *S*'s *knowledge* that *P* is infallible iff it's impossible that *S* believes that *P* on the same basis and it's false that *P* (= C_K).

A belief, justification or piece of knowledge is *fallible* iff the respective condition isn't met. "Justification" is used as a placeholder here and should be understood to cover internalist and externalist justification, entitlement, warrant and so on. Note that the definitions of infallibility aren't entirely parallel. Obviously, it's impossible to know that *P* when it's false that *P*. But from this it doesn't follow that knowledge is trivially infallible. Hence, infallibility of knowledge should be understood as the impossibility to believe on the same basis (or with the same justification) while it's false that *P*.

A more difficult problem is to decide whether belief, justification or knowledge should be the proper object of the debate between fallibilism and infallibilism. Since knowledge that *P* is infallible just in case the justification for the belief that *P* is infallible, (in-)fallibilism about justification and knowledge are closely connected. Since I'm interested in the relation between (in-)fallibilism and scepticism and have formulated scepticism as a thesis about knowledge, I will restrict attention to (in-)fallibilism about knowledge.[6] But what about belief? Obviously, neither are all beliefs infallible nor is it a necessary condition for having a belief that it is infallible. Although the *distinction* between fallible and infallible beliefs is a sensible one, fallibilism and infallibilism as *theses* are only interesting when applied to justification or knowledge. When fallibilism is formulated as a thesis about belief, we're usually dealing with the thesis that no belief – no matter whether and how well it's justified – is fallible. Fallibilism about belief in this sense is a consequence of fallibilism about justification and the debate should focus on the latter claim.

This observation leads to the second problem, the need to distinguish between conditional and unconditional versions of fallibilism and infallibilism. *Conditional* (in-)fallibilism consists in the acceptance or rejection of a necessary condition for knowledge, whereas unconditional (in-)fallibilism consists in a thesis about whether the condition is ever met or not.

[6] A second reason for focussing on (in-)fallibilism about knowledge is that conditional infallibilism (see below) is more plausible when understood as a thesis about knowledge – knowledge requires infallibility – than when understood as a thesis about justification – justification requires infallibility. An infallibilist about knowledge can concede that e.g. inductive beliefs are justified while maintaining that they never amount to knowledge since they aren't infallible.

> *Conditional infallibilism about knowledge:* C_K is a necessary condition for knowing that *P*, i.e.: If *S* knows that *P*, it's impossible that *S* believes that *P* on the same basis and it's false that *P*.

Accordingly, conditional fallibilism about knowledge is defined by rejecting C_K as a necessary condition for knowledge.

Unconditional infallibilism is defined by the following thesis about the human condition:

> *Unconditional infallibilism about knowledge:* At least sometimes C_K is met, i.e. sometimes it's impossible that *S* believes that *P* on the same basis and it's false that *P*.

Accordingly, unconditional fallibilism about knowledge is defined by the negation of this thesis: C_K is never met, i.e. it's always possible that *S* believes that *P* on the same basis and it's false that *P*.

Conditional and unconditional infallibilism are independent of each other. One can require of knowledge that it must be infallible and combine this with either the claim that no knowledge claim meets this condition or with the claim that some knowledge claims do meet it. This results in the four options that are depicted in table 1.

	C_K necessary for *K*	C_K not necessary for *K*
C_K sometimes met	Cartesian infallibilism	(optimistic fallibilism)
C_K never met	sceptical infallibilism	Peircean fallibilism

Table 1**:** Options

All the positions from the table are conjunctions of two claims. Cartesian infallibilism asserts that (a) knowledge requires infallibility and (b) that some of our claims to knowledge are infallible. Peircean fallibilism[7] asserts that (a) knowledge doesn't require infallibility and (b) that all of our knowledge claims are fallible. Sceptical infallibilism partially agrees with both of them. It agrees with the first claim of Cartesian infallibilism and with the second claim of Peircean fallibilism. (Optimistic fallibilism is a position I mention mainly for the sake of comprehensiveness. It agrees with Cartesian infalli-

[7] Peirce is often credited with introducing the term "fallibilism" in this sense. However, it's not beyond dispute what sort of fallibilism should be called "Peircean fallibilism". Rott e.g. uses "broadly Peircean or Popperian notion of fallibilism" to stand for what is called "unconditional fallibilism" here and even for the stronger thesis that "for all subjects and at all times, some of the subject's beliefs are wrong" (Rott 2005, p. 475 and n. 14).

bilism that some of our knowledge claims are infallible, but holds that this doesn't really matter since infallibility isn't necessary for knowledge.)

Two lessons are to be learnt from this terminological exercise. The first lesson is to keep in mind that two issues are at stake when discussing fallibilism and infallibilism. One issue is whether a particular condition is necessary for knowledge, the other issue is whether this condition is ever met. The second lesson is to keep in mind the entailments between the various positions and scepticism. As defined above, Cartesian infallibilism entails that we do know something whereas sceptical infallibilism entails that we don't know anything. Peircean fallibilism – again: as defined above – entails neither that empirical knowledge is possible nor that it's impossible. That is so because the position says nothing about what the true necessary conditions for knowledge are; it only claims that infallibility isn't among them. Among the true necessary conditions there still might be a condition that it's impossible to meet. This is a departure from common usage; usually Peircean fallibilism is understood to be committed to the claim that we do know something or that knowledge is at least possible. For my purposes this departure is useful, however. I will argue below that even the Peircean fallibilist should accept a necessary condition for knowledge different from the infallibility requirement that we can't fulfil.

2.2.2. *The Infallibility Condition*

In the discussion so far I have left open which kind of modality is meant in the above definitions. Here are some candidates for the infallibility condition (cf. Dutant 2007):

Logical Impossibility: It's logically impossible to believe that *P* (on basis *B*) and *not-P*.

Epistemic Impossibility: It's epistemically impossible to believe that *P* (on basis *B*) and *not-P*.

(a) *Ruling Out Condition: S* can rule out all error-possibilities (i.e. possible worlds in which *P* is false).

(b) *Evidence Condition: S* has some evidence *E* for believing *P* and *E* guarantees the truth of *P* (or in terms of epistemic probability: $Pr(P|E)=1$).

Metaphysical Impossibility: It's metaphysically impossible to believe that *P* (on basis *B*) and *not-P*.

The parenthetical "on basis *B*" is needed to deal with a problem akin to Nozick's grandmother. A belief isn't fallible just because it might be false when based on a fancy unreliable method not actually relied upon. It's infallible if it's impossibly false when based on the same method actually relied upon.

The first proposal – logical impossibility – is obviously far too strong (cf. Dutant 2007, p. 59). Even Descartes doesn't claim logical infallibility for knowledge. The second proposal – epistemic impossibility – needs a closer look, but is, as I will argue, ultimately unconvincing. When explaining infallibility in terms of epistemic necessity, two proposals stand out. The idea behind the ruling out condition is that a belief is infallible if the believer is able to rule out all possible worlds in which *P* is false. The idea behind the evidence condition is that a belief is infallible if the believer's evidence is truth-guaranteeing. To evaluate the ruling out condition let us first clarify what "ruling out" means. The most natural reading of "ruling out" is that *S* can rule out an error-possibility iff *S* knows, or is in a position to know, that this error-possibility doesn't obtain. But then every conception of knowledge that secures closure is classified as infallibilist. Take Sartwell's definition of knowledge as merely true belief as an example (Sartwell 1992; this argument is taken from Dutant 2007, p. 66). If I believe truly that I have hands and believe, or am at least in a position to believe, that this implies that I'm not a brain in a vat, I believe truly and, hence, know that I'm not a brain in a vat. But true belief is obviously not infallible in any sense worth that name! Hence, the ruling out condition captures the core of infallibilism only if more is read into "ruling out" (cf. Dutant 2005, pp. 62-66). The evidence condition is more specific than the ruling out condition. The evidence condition implies the ruling out condition, but not *vice versa*. However, the evidence condition suffers from the same defect as the ruling out condition. McDowell's disjunctivism and Williamson's *E=K*-ism are views that are classified as infallibilist by this definition, but are not intuitively infallibilist. According to both, in a possible world in which *P* is false *S* (in many cases) wouldn't have the same evidence. Hence, one's evidence guarantees the truth of *P*. However, this result depends on McDowell's and Williamson's externalist construal of evidence. According to McDowell (1982, 1995), my evidence can be something like *I see a red apple* even if I can't distinguish seeing a red apple from hallucinating a red apple. According to Williamson (2000, ch. 9), my evidence consists in everything I know: if I know that there is a red apple, that is part of my evidence. Hence, my evidence guarantees that there is a red apple. Whatever the merits of such positions, they should surely not count as infallibilist.

Unfortunately, all epistemic conditions face the same problem: the epistemic terms ("ruling out", "evidence", "know" etc) used to define fallibilism can themselves be understood in a fallibilist and an infallibilist way. If "ruling out" means "can know not to obtain", the ruling out condition can only be used to define fallibilism if knowledge is already understood in an infallibilist sense. If my evidence is equated with everything I know, the evidence condition can, again, only be used to define fallibilism, if knowl-

edge is already understood in an infallibilist sense. The problem is a principled one: definitions of infallibilism in epistemic terms are only adequate if the respective epistemic terms are already understood in an infallibilist way. I propose to circumvent this problem by not using epistemic terms in formulating the infallibility condition at all.

According to the metaphysical impossibility condition, a belief on basis *B* is infallible iff it could not have been false when based on *B*.[8] This condition almost avoids the problem of using epistemic terms. The "almost" is due to the "on basis *B*" still present in the condition. Suppose we understand that clause in such a way that e.g. the basis of my belief that there's a red apple is my seeing a red apple. Then it's metaphysically impossible to have that belief on that basis when it's false. In a possible world in which there's no red apple I can't *see* that there's a red apple. To avoid this problem I need to replace the clause on basis *B* by a clause that is neutral with respect to different understandings of what epistemic bases are. I propose a doppelgänger condition as a working definition. A doppelgänger is an internal duplicate. Internal duplicates share mental states in the narrow sense. Such a doppelgänger is in the same brain state, but may or may not be in the same environment.

> *The Doppelgänger Condition: S*'s belief that *P* is infallible iff the corresponding belief of every doppelgänger of *S* is true.

According to this condition, a belief is infallible iff it could not have been false. The clause "corresponding belief" solves a problem with the "it" in "*it* could not have been false". The doppelgänger condition shouldn't assume that beliefs have to be individuated internalistically. If they aren't, not all of my doppelgängers have the same beliefs I have. Suppose that some of my doppelgängers don't believe that water quenches thirst because they are in an XYZ environment. What corresponds to my belief that water quenches thirst is their belief that XYZ quenches thirst. My belief is infallible only if their corresponding beliefs are true as well.

With this infallibility condition in place it's useful to introduce a new definition of *ruling-out*. This definition helps to avoid cumbersome repetitions of the doppelgänger condition.

[8] Note that this condition is supposed to range over *all* possible worlds. If "could" ranges over nearby possible worlds only, some beliefs are infallible just because they are false only in remote possible worlds. If such a restriction were in place, any safety account of knowledge – *S* knows that *P* iff *S*'s belief could not easily have been false – would turn out to be an infallibilist conception of knowledge. However, intuitively safety accounts aren't infallibilist accounts. To avoid this problem we need to stipulate that all possible worlds are being quantified over.

Ruling-out: S can rule out an error-possibility for the belief that *P* iff there's no doppelgänger of *S* with a corresponding false belief.

In what follows "ruling-out" is always to be understood in this sense.

The resulting definition of conditional fallibilism is a demanding one: one can know something only if no doppelgänger believes it falsely. Probably, only Descartes and some other rationalists are infallibilists in this sense. But this is fine for my purposes. The challenge to be answered below is that the sceptical argument rests on infallibilism and that infallibilism is too strong a requirement for knowledge. This challenge would be implausible on a relaxed definition of infallibilism that included e.g. McDowell and Williamson as infallibilists. Since their positions are designed to avoid scepticism, the sense of "infallibilism" on which the challenge is based cannot be too relaxed – otherwise the challenge would be unconvincing from the very beginning.

2.3. *Summary*

In this section I have outlined scepticism, fallibilism and infallibilism. In the next two sections I argue that scepticism is to be distinguished from both fallibilism and infallibilism. The distinctions drawn in this section offer a fresh view on the relations. To sum up, it seems that scepticism agrees partly with both fallibilism and infallibilism. On the one hand, scepticism *agrees* with unconditional fallibilism that every empirical belief is fallible. The question to ask is where they *dis*agree. Is the disagreement between sceptics and fallibilists only concerned with the definition of knowledge, hence a merely verbal disagreement? On the other hand, scepticism agrees with conditional infallibilism that infallibility is a requirement for knowledge: only if infallibility is required for knowledge, can the sceptic infer the impossibility of empirical knowledge from the fallibility of every empirical belief. However, if this is so, the force of the sceptical argument depends on the plausibility of conditional infallibilism. In the next two sections I reassess both claims about the relation between scepticism and fallibilism/infallibilism.

3. *Scepticism vs Fallibilism*

The task for this section is to show that scepticism and fallibilism are not just the same view (i.e. every empirical belief is fallible) under different guises. *Prima facie* the respective claims of fallibilism and scepticism are as follows:

The Fallibilist: There is a doppelgänger of me who believes falsely that he has a hand, yet I know whether I have hands.

The Sceptic: There is a doppelgänger of me who believes falsely that he has a hand, hence I don't know whether I have hands.

Understood this way, they agree about some factual matters, but disagree about what follows from them. Put more precisely, they agree about all propositions not containing the concept *knowledge*, i.e. they agree about all the basic facts on which knowledge-facts supervene. If we eliminated the word "knowledge" and its cognates from our language, all disagreement would vanish. Hence, the dispute between the fallibilist and the sceptic seems to be a merely verbal dispute. This line of reasoning results in a challenge for scepticism: is there a substantial disagreement between scepticism and fallibilism?[9]

3.1. *A Scope Distinction*

Despite appearances there is a substantial difference between fallibilism and scepticism: it can be explained as a scope distinction. It is well-known that sentences like "at some time everything ends" are ambiguous: *either* everything will end at the same time *or* everything will end sometime, but possibly at different times. The same ambiguity applies to "every empirical belief is fallible":[10]

Fallibilism: (F) For every of *S*'s empirical beliefs there's a possible world in which *S*'s doppelgänger's *belief* is false.

Scepticism: (S) There's a possible world in which all of *S*'s doppelgänger's empirical beliefs are false.

Anti-Sceptical Fallibilism: Fallibilism, but not scepticism, i.e. (F), but not (S).

(S) is stronger than (F). While (S) entails (F), (F) doesn't entail (S). If (S) is true, there is a global error-possibility, i.e. a possibility in which all empirical beliefs are false at the same time. For (F) to be true it suffices that there is a local error-possibility for each empirical belief, i.e. a possibility in which the target belief is false while all other beliefs may be true.

[9] In this paper I won't deal with the reply that the substantial disagreement is a normative one. For example, the sceptic and the fallibilist might disagree about whether it's permissible to assert that I have hands. Nothing I say here is incompatible with that reply.

[10] Of course, I'm not the first one to notice this scope distinction. As Wittgenstein remarks in *On Certainty*, that nothing is exempted from doubt doesn't mean that everything can be doubted at once, but only that anything can be doubted (1969, §115, §232). This distinction is also drawn by Davidson (1983, p. 140).

While there's clearly a scope ambiguity in "every belief could be false", it's not obvious that this ambiguity marks the distinction between scepticism and fallibilism. The basic idea behind these definitions is that anti-sceptical fallibilism sets a limit to our fallibility: while it admits that each empirical belief could be false, it keeps up hopes by maintaining that not all empirical beliefs could turn out to be false. If that was the case, it would amount to some sort of damage control: it's not troublesome that I can't rule out some *local* error-possibilities because even if they should obtain, only some of my beliefs would be affected. Nevertheless, it's troublesome that I can't rule out *global* error-possibilities. If such a possibility obtained, my complete system of empirical belief would be affected. What turns fallibilism into an anti-sceptical stance, is some assurance that even if no empirical belief is guaranteed to be true, there's at least some reason to think that the worst possibility – global error – can't obtain. According to scepticism, although fallibilism has learned a lesson, it stops short of the true lesson. Moreover, if the sceptic only claimed that each belief is fallible, she could present a different error-possibility for every empirical belief. I don't know that the person across the street is Anne because it might be her identical twin nobody has ever heard of. I don't know that it's 2013 because it might be 2014 and the whole world was drugged for a year. There should be a reason why the sceptic doesn't come up with a new error-possibility for every belief, but chooses to stick with a single scenario like the brain in a vat scenario. She isn't interested in piling up error-possibilities – or so my diagnosis goes – because to these challenges one could indeed reply that the existence of specific error-possibilities only proves fallibilism.

3.2. *Error-Possibilities or Ignorance-Possibilities?*

Unfortunately, (S) is demonstrably false: Not all of my empirical beliefs can be false at the same time.[11] Among my empirical beliefs are for example:

(7) There are animals.

(8) There are no unicorns.

But if (7) is false, (8) must be true and *vice versa*. Error-possibilities for (7) are worlds without animals, error-possibilities for (8) are worlds that contain at least one animal, namely a unicorn. No matter how severe my epistemic misfortune I will always be fortunate enough to have some true beliefs. This isn't just a quirk. A consistent set of beliefs (of sufficient complexity) is usually

[11] As far as I know, the earliest discussion of this point is Müller (2003, p. 47). It's also discussed in Gemes (2009, 2010), Genova (2010), LittleJohn (2012, p. 131) and Kraft (2013).

not falsity-consistent, i.e. a set of propositions is often consistent even if the set of the negations of those propositions isn't consistent. Hence, even though each of my empirical beliefs could be false, they can't all be false at the same time.

Even though global error is impossible, the sceptic doesn't have to give up yet. If sceptical scenarios don't illustrate the possibility of global error, they should be understood as achieving something different. Consider how a sceptic would argue that a brain in a vat doesn't know that it has a brain, that it has less than seven hands, that computers exist or that all firefighters are courageous. All of these beliefs are true, even in the brain in a vat scenario. Yet, a brain in a vat can know none of those true propositions. I think that it's more or less obvious that a brain in a vat doesn't know that it has a brain and so on. In case an argument is needed there are different paths to reach that conclusion. One path relies on the thesis that an accidentally true belief can't be knowledge: the brain in a vat's true beliefs are cases of veritic epistemic luck in Pritchard's terminology (2005, p. 146).[12] This claim doesn't depend on either an internalist or an externalist conception of knowledge: for neither the evidence of an envatted brain nor its belief-forming process leading to the belief that she has a brain has anything to do with her actual brain. The evidence is concerned with simulated humans and their anatomy. The belief-forming process is triggered by stimulations received from the supercomputer. Hence, the epistemic basis of the belief has nothing to do with the actual brain. This is also the case for the beliefs that there are no unicorns or that all firefighters are courageous. The epistemic basis for the first belief has to do with the absence of unicorns in the simulation, not with their absence around the vat. The epistemic basis of the second belief has to do with the simulated behaviour of simulated firefighters and not with their non-existence. If the epistemic basis of a true belief is in that way unconnected to the fact, the belief is only accidentally true and, therefore, it doesn't amount to knowledge.

Hence, the most charitable interpretation of sceptical scenarios understands them as illustrating the possibility of global ignorance, the possibility of not *knowing* anything empirically. A sceptical scenario is not defective just because it does not illustrate the possibility of global error. To be successful it has to illustrate the possibility of global empirical ignorance, i.e. the possibility that all of my empirical beliefs fall short of knowledge. If sceptical scenarios are ignorance-possibilities, the untenability of (S) shouldn't worry the sceptic. (S) needs to be revised, but that revision does no harm:

[12] Another path, which isn't developed here, is to rely on a causal requirement for knowledge.

(S*) There's a possible world in which none of *S*'s doppelgänger's empirical beliefs amounts to knowledge.

(F) should be revised accordingly:

(F*) For every of *S*'s empirical beliefs there's a possible world in which *S*'s doppelgänger's belief doesn't amount to knowledge.

With these revisions my original suggestions can be upheld. Scepticism is stronger than fallibilism because the latter is already true if for each empirical belief there's an ignorance-possibility whereas the former is true only if there's a single ignorance-possibility for all empirical beliefs. To sum up, my proposal for distinguishing scepticism and fallibilism has two notable consequences. First, the sceptical scenario has to be understood as an ignorance-possibility. Second, this distinction requires that the sceptical argument work with a single, global sceptical scenario. The scenario may not be adjusted on a case by case basis, i.e. the argument may not offer the scenario of a handless brain in a vat to attack the belief that I have hands and the scenario of a seven-handed envatted creature to attack the belief that I have less than seven hands.

4. *Scepticism vs Infallibilism*

The substantial issue between (anti-sceptical) fallibilism and scepticism should be whether there is a global ignorance-possibility for all empirical beliefs or only local ignorance-possibilities for each empirical belief. Moreover, I have argued that at least one such global ignorance-possibility exists. The original brain in a vat scenario in which the whole universe consists of only the envatted brain and the supercomputer both created by chance at the beginning of the universe does the trick. Such a brain in a vat doesn't know that it has a brain or that computers exist because its sense experiences aren't a means to find out anything about the external world, but at most a means to find out some things about the simulation.

What I haven't argued so far is that this result helps to reassess the relation between scepticism and infallibilism. To defend scepticism what is needed is not only a global ignorance-possibility, but also (a) a reason why we *can't* know that this ignorance-possibility doesn't obtain (i.e. the first premise of the argument from ignorance), and (b) a reason why we *need* to know that this ignorance-possibility doesn't obtain to know anything empirically (i.e. the second premise of the argument from ignorance). *Prima facie*, conditional infallibilism is needed for both tasks. If I have to rule out *all* ignorance-possibilities, *a fortiori* I have to rule out the brain in a vat scenario. But I don't know that it doesn't obtain because – according to conditional

infallibilism – I know that only if all of my doppelgängers know that they aren't a brain in a vat. Of course, this only shows that the sceptical argument can be defended with the help of conditional infallibilism; it doesn't show that this is the only possible defence. Nevertheless, Brueckner has argued recently that scepticism is tied to infallibilism:[13] without infallibilism there's no reason to think that I need to, but can't know whether I'm a brain in a vat. In this section I argue that the distinction between local and global ignorance-possibilities helps to answer Brueckner's challenge.[14]

4.1. *The Relevance of Global Ignorance-Possibilities*

Standard analyses of the sceptical argument assume the principle of unrestricted relevance. According to this principle, any error/ignorance-possibility whatsoever is relevant just because it's an error/ignorance-possibility. This principle allegedly solves the problem that intuitively the sceptical scenario isn't more relevant than, say, some crazy conspiracy theory involving secret agents from another planet.[15] That is a mistake: to defend scepticism it should be explained, not just assumed why sceptical scenarios are supposed to be relevant. In what follows I argue that there is a feature of sceptical scenarios that makes them relevant. If that idea can be put to work, scepticism can be severed from infallibilism: scepticism is compatible with the intuitive idea that one's knowledge that, say, this is a zebra isn't threatened just because one can't rule out that it's a painted mule. Thereby, a sceptic can admit that not all ignorance-possibilities need to be ruled out, i.e. that some ignorance-possibilities are irrelevant or too far-fetched or not worthy our time. Globality is a special feature that distinguishes genuinely threatening ignorance-possibilities from other ignorance-possibilities (painted mules, conspiracy theories, identical twins nobody knows of etc). However, in order to formulate a principle on behalf of the sceptic, I need to clarify

[13] According to Brueckner, "it turns out that a canonical form of skeptical argument depends upon the denial of fallibilism" (2005, p. 384). In a later paper, his conclusion is a bit more tentative: "This paper seems to end not with a bang but a whimper on the skeptic's part. It's just not entirely clear whether the charge that the skeptical argument collapses into Infallibilism can be successfully answered by the skeptic" (2011, p. 87).

[14] For other replies to Brueckner, see Briesen (2010) and Dodd (2012).

[15] A nice example for this understanding of scepticism is Lewis (1996). Lewis assumes without further argument that the sceptical context is a context in which no error-possibility may be properly ignored. I think that even according to the sceptic some error/ignorance-possibilities are legitimately ignored. What is characteristic of the sceptic is the rule of globality: no *global* error/ignorance-possibility may be properly ignored. Thus the sceptic need not be saddled with the over-demanding claim that even conspiracy theories can't be properly ignored.

what globality amounts to. Of course, sceptical scenarios aren't global in the sense that in them all beliefs whatsoever don't amount to knowledge. They are only global in the sense that all beliefs based on a particular evidential source or acquired via a particular belief-forming process are affected. I will formulate the needed principle in terms of capacities to stay neutral between internalist and externalist conceptions of epistemic bases (evidential sources, belief-forming processes etc). Since a belief can be based on more than one evidential source (e.g. perception and inference), an ignorance-possibility for a capacity affects all beliefs based on that capacity and, possibly, other capacities. For example, the brain in a vat scenario is a global ignorance-possibility with respect to all beliefs based inter alia on perception, the scenario of the recently created earth is a global ignorance-possibility with respect to all beliefs based inter alia on memory and so on. Hence, I propose the following principle on behalf of the sceptic:

> *Globality Principle: S* knows that *P* via capacity *C* only if *S* knows of any *C*-global ignorance-possibility that it doesn't obtain.

This principle explains how the sceptic can maintain that I *need* to rule out being a brain in a vat without maintaining that I need to rule out any ignorance-possibility whatsoever. But more needs to be said on why a Peircean fallibilist (who accepts conditional and unconditional fallibilism) should accept the globality principle. Moreover, more needs to be said on why I *can't* know not being a brain in a vat. The next two sections deal with these two tasks.

4.2. *Fallibilism and the Globality Principle*

So far I haven't discussed *why* fallibilists think that knowledge is compatible with the possibility of being wrong. While some fallibilists seem to rest their case merely on intuitions about examples, some give more explicit arguments. According to the first strategy, it's intuitively true that one can know whether this is a zebra without being able to rule out that it's a painted mule. The latter possibility is just too remote to be relevant. Consideration of such cases, however, doesn't help to find out whether the brain in a vat scenario is relevant or not. While some think that this possibility is obviously too remote, others can't help but think that it's too important to be dismissed that easily. Indeed, the sceptic can concede that the painted mule possibility isn't relevant while maintaining that the brain in a vat possibility is relevant, i.e. that relevance isn't correlated with closeness. The case-based argument for fallibilism only shifts the burden of proof: the sceptic has to argue that the brain in a vat possibility is relevant. Whether she can succeed

with that task, just can't be decided by relying on the intuitive case for fallibilism.

What about the second strategy, offering a general argument for the compatibility of knowledge with the possibility of being wrong? Such an argument is at least hinted at by e.g. Peirce and Popper. According to Peirce, it's important for our claims to knowledge that we as inquirers can get closer to the truth in the long run. According to Popper, we might not get closer to the truth in the long run, but can at least falsify our hypotheses about the external world while hypotheses we can't falsify are rejected as bogus. The common idea here is that although every belief can be false, the false ones will or at least can get sorted out. We can be satisfied with fallible knowledge as long as piecemeal improvement of our system of empirical belief is likely or at least possible. They hope that sooner or later false beliefs will be revealed as such: even if none of our empirical belief is safe from being false, we can get closer to the truth by adjusting our beliefs to the ever increasing body of evidence, re-applying our epistemic methods and so on.

This line of reasoning is an argument for conditional fallibilism because it shows why we need not worry too much about the possibility of being wrong. Nevertheless, this argument offers the sceptic a point of attack: it's convincing only if it's indeed plausible that it's at least possible that our belief system is improved by piecemeal adjustment. However, as long as I can't rule out being a brain in a vat, I have no reason to think that piecemeal adjustment does lead to an improved belief system. A brain in a vat revises her empirical beliefs all the time, but this doesn't bring her closer to the truth. No matter how careful and rational it is in revising her empirical beliefs, the result won't be knowledge. Suppose for the sake of the argument that I could rule out being a brain in a vat. Then the above line of reasoning offers a good reason for believing something even if I presently can't rule out all ignorance-possibilities. If I'm not a brain in a vat, it's at least possible that my false beliefs will be falsified by future evidence and replaced by better beliefs. Unless the condition imposed by the globality principle is met, there's no reason to think that piecemeal adjustment of our beliefs will indeed result in an improvement of our beliefs. The globality principle just spells out a requirement for the possibility of piecemeal improvement. Hence, this principle isn't a condition imposed by the sceptic; it formulates a precondition of piecemeal improvement of our beliefs heralded by the fallibilist.

One might object that the possibility of piecemeal improvement is secured if *as a matter of fact* I'm not a brain in a vat – no matter whether I *know* it or not. After all I'm better off if I have some empirical beliefs than if I have no empirical beliefs at all. However, the salient alternative to my actual beliefs is not having no beliefs whatsoever, but believing that I'm a brain in

a vat. This would lead to a widespread, but easy to accomplish adjustment of my empirical beliefs. Instead of believing that I see a red apple I would believe that I see a simulated red apple, instead of believing that I have two hands I would believe that I have two simulated hands, instead of believing that there are animals I would believe that there are animals in the simulation, instead of believing that relying on my sense experiences is a reliable method of forming beliefs about the external world I would believe that this is a reliable method for forming beliefs about the simulation, and so on. We're not coerced into one or the other system of beliefs. We're free to believe that we're brains in a vat and hedge all empirical beliefs with an "in the simulation"-operator. Hence, if the fallibilist claims that I don't need to know whether a global ignorance-possibility like the brain in a vat scenario obtains, I'm left free to switch to the alternative system of empirical beliefs. Of course, the switch wouldn't be backed up by any reasons, but that only means that it would be arational, not that it would be irrational. Of course, given that I believe that I'm not a brain in a vat, it would be irrational and unjustified for me to believe that I see a simulated red apple and so on. That is, once I've decided to believe that I'm not a brain in a vat, my further beliefs aren't subject to arbitrary decisions. As soon as I become aware of the brain in a vat scenario, I can decide whether I believe to be one or not to be one. Moreover, with respect to some local ignorance-possibility I can decline to make up my mind without risking irrationality and withhold belief, but with respect to global ignorance-possibilities that isn't an option (unless I'm willing and able to withhold all empirical beliefs). Since the decision affects all my empirical beliefs, the fallibilist shouldn't claim that global ignorance-possibilities are irrelevant and that I don't need to know whether they obtain.

4.3. *Fallibilism and the Non-Arbitrariness Principle*

To complete the argument I need to show why fallibilists should agree that I don't know whether I'm a brain in a vat. Given the many accounts of knowledge defended in the literature that's a huge task, but I will nevertheless offer a suggestion. Let us consider three common defences of the claim that I don't know whether I'm a brain in a vat. The first prominent defence relies on sensitivity: if I were a brain in a vat, I would still believe that I'm not brain in a vat. Hence, this belief is insensitive and doesn't amount to knowledge. A second prominent defence invokes the sameness of evidence lemma: a brain in a vat has the same evidence as I have. Therefore, my evidence doesn't justify believing being a brain in a vat. Both these defences rely on claims about knowledge/justification which are too contentious for present purposes (and the latter one leads to infallibilism). A third common

defence argues that my justification for believing that I'm not a brain in a vat is circular: even if my evidence includes that I see my hand or that I have a hand, this evidence can't be used as evidence for believing that I'm not a brain in vat; that would be circular. Since the brain in a vat scenario affects all empirical beliefs my justification is circular without an independent justification for trusting my senses.[16] I want to offer another defence which is akin to the circularity defence but avoids the difficulty of clarifying the relevant notion of circularity; it's implicitly contained in the last paragraph. I'm free to believe that I *am* a brain in a vat and to adjust my beliefs accordingly. A convincing anti-sceptical view should not just explain why I am justified in believing that I'm not a brain in a vat, but also why I would not be justified in believing that I'm a brain in a vat if I should choose to believe so. If both claims are equally justified, it's arbitrary which of the two belief systems I accept and, hence, I don't know whether I'm a brain in a vat.

> *Non-Arbitrariness Principle:* Suppose *not-P* is a *C*-global ignorance-possibility. Then I know that *Q* via *C* only if it isn't epistemically arbitrary whether I believe that *P* or that *not-P* (assuming I would make the appropriate adjustments in my belief system).[17]

What is compared here are two beliefs (or two belief systems), not a belief with withholding that belief. Admittedly, I have some justification for believing that I'm not a brain in a vat, maybe in the form of a default justification (Williams 2001) or an entitlement (Wright 2004). Hence, that belief might well be better justified than withholding it. What I argue here is that such a justification only leads to knowledge if my belief that I'm not a brain in a vat is better justified than its rival belief would be. So does the belief that I'm not a brain in a vat meet the non-arbitrariness requirement? Since my aim isn't to defend scepticism, but merely to present a cogent sceptical argument that doesn't assume infallibilism, and since it would be impossible here to go through all anti-sceptical strategies, I rest content with posing a challenge: a plausible anti-sceptical strategy should show how the non-arbitrariness requirement can be met. At least, it's difficult to see how it could be met. If my justification for believing that I'm not a brain in a vat depends on other empirical beliefs, my doppelgänger who is exactly like me but believes that he's a brain in a vat and adjusts his other beliefs accordingly is

[16] Dodd offers a sceptical argument freed from infallibilism by relying on such a circularity argument (2012, pp. 349-351).

[17] The parenthesis needs some explanation. Which adjustments are appropriate depends on the ignorance-possibility considered. If I believed the ignorance-possibility that I'm a brain in a vat, I would need to adjust every empirical belief, e.g. by hedging it with the "in the simulation"-operator.

justified to the same degree. If I'm entitled to my belief without possessing evidence for it, my doppelgänger enjoys the very same entitlement: his belief is as much a cornerstone or precondition for his cognitive projects as my belief is. If I may dogmatically rely on my perceptual beliefs (*there's an apple*), my doppelgänger (*there's a simulated apple*) may do so, too. Hence, it's not far-fetched to presume that my belief that I'm not a brain in a vat is arbitrary or, to vary Russell's *poached egg*-dictum, that someone who believes to be a brain in a vat is only in the minority.

5. *Summary*

The aim of this paper was to reassess the relation between scepticism on the one hand and infallibilism and fallibilism on the other hand. The common theme for that reassessment was the idea that sceptical scenarios are global ignorance-possibilities. Based on that idea I argued that scepticism is stronger than and independent of infallibilism. To achieve the latter goal I appealed to two principles, the globality and the non-arbitrariness principles. The globality principle explains why global ignorance-possibilities are relevant ignorance-possibilities. The non-arbitrariness principle explains why I can't know whether a global ignorance-possibility obtains. I have offered some reasons why both principles should be accepted, even by fallibilists. Whether that means that scepticism is ultimately true, I'd rather not say. Needless to say, I haven't offered conclusive reasons against the possibility of a version of fallibilism that rejects at least one of the two principles. Nevertheless, they pose a challenge for fallibilist attempts at solving the sceptical problem.[*]

Department of Philosophy
University of Regensburg
Germany

E-mail: tim.kraft@psk.uni-regensburg.de

[*] I'm grateful to the audience of my talk at the European Epistemology Network Meeting at Bologna and Modena and to Stefan Ruhland, Hans Rott, Giorgio Volpe and two anonymous referees for helpful comments on earlier versions.

References

Briesen, J. 2010: "Reconsidering Closure, Underdetermination, and Infallibilism", *Grazer Philosophische Studien*, 80, pp. 221-236.

Brueckner, A. 2005: "Fallibilism, Underdetermination, and Skepticism", *Philosophy and Phenomenological Research*, 71, pp. 384-391.

Brueckner, A. 2011: "*~K~SK*", *Philosophical Issues*, 21, pp. 74-89.

Davidson, D. 1983: "A Coherence Theory of Truth and Knowledge"; page references to reprinting in Id., *Subjective, Intersubjective, Objective*, Oxford, Clarendon, 2001, pp. 137-153.

DeRose, K. 1995: "Solving the Skeptical Problem", *Philosophical Review*, 104, pp. 1-52.

Dodd, D. 2011: "Against Fallibilism", *Australasian Journal of Philosophy*, 89, pp. 665-685.

Dodd, D. 2012: "Evidentialism and Skeptical Arguments", *Synthese*, 189, pp. 337-352.

Dougherty, T. 2011: "Fallibilism", in S. Bernecker and D. Pritchard (eds.), *The Routledge Companion to Epistemology*, London, Routledge, pp. 131-143.

Dutant, J. 2007: "The Case for Infallibilism", in C. Penco, M. Vignolo, V. Ottonelli and M.C. Amoretti (eds.), *Proceedings of the 4th Latin Meeting in Analytic Philosophy, LMAP/07*, Genoa (Italy), September 20-22, 2007, CEUR Workshop Proceedings, pp. 59-84, http://ceur-ws.org/Vol-278/paper05.pdf.

Gemes, K. 2009: "A Refutation of Global Scepticism", *Analysis*, 69, pp. 218-219.

Gemes, K. 2010: "A Vindication of a Refutation of Global Scepticism, a Refutation of Global Perceptual Scepticism and a Refutation of Global Existential Scepticism", *Analysis*, 70, pp. 63-70.

Genova, A.C. 2010: "Has Gemes Refuted Global Scepticism?", *Analysis*, 70, pp. 59-63.

Hetherington, S. 2012: "The Gettier-Illusion: Gettier-Partialism and Infallibilism", *Synthese*, 188, pp. 217-230.

Howard-Snyder, D., Howard-Snyder, F. and Feit, N. 2003: "Infallibilism and Gettier's Legacy", *Philosophy and Phenomenological Research*, 66, pp. 304-327.

Klein, P. 1981: *Certainy. A Refutation of Scepticism*, Minneapolis, University of Minnesota Press.

Kraft, T. 2013: "Sceptical Scenarios are not Error-Possibilities", *Erkenntnis*, 78, pp. 59-72.

Leite, A. 2010: "Fallibilism", in J. Dancy, E. Sosa and M. Steup (eds.), *Blackwell Companion to Epistemology*, 2nd ed., Malden, Blackwell, pp. 370-375.

Lewis, D. 1996: "Elusive Knowledge"; page references to reprinting in Id., *Papers on Metaphysics and Epistemology*, Cambridge, Cambridge University Press, 1999, pp. 418-445.

LittleJohn, C. 2012: *Justification and the Truth-Connection*, Cambridge, Cambridge University Press.

McDowell, J. 1982: "Criteria, Defeasibility, and Knowledge", *Proceedings of the British Academy*, 68, pp. 455-479.
McDowell, J. 1995: "Knowledge and the Internal", *Philosophy and Phenomenological Research*, 55, pp. 877-893.
Moon, A. 2012: "Warrant Does Entail Truth", *Synthese*, 184, pp. 287-297.
Müller, O. 2003: *Wirklichkeit ohne Illusionen I. Hilary Putnam und der Abschied vom Skeptizismus oder: Warum die Welt keine Computersimulation sein kann*, Paderborn, Mentis.
Pritchard, D. 2005: *Epistemic Luck*, Oxford, Clarendon.
Reed, B. 2002: "How to Think About Fallibilism", *Philosophical Studies*, 107, pp. 143-157.
Rott, H. 2005: "Stability, Strength and Sensitivity: Converting Belief Into Knowledge", *Erkenntnis*, 61, pp. 469-493.
Sartwell, C. 1992: "Why Knowledge Is Merely True Belief", *The Journal of Philosophy*, 89, pp. 167-180.
Stroud, B. 1984: *The Significance of Philosophical Scepticism*, Oxford, Clarendon.
Unger, P. 1975: *Ignorance. A Case for Scepticism*, Oxford, Clarendon.
Wang, L. and Tai, O. 2010: "Skeptical Conclusions", *Erkenntnis*, 72, pp. 177-204.
Williams, M. 1991: *Unnatural Doubts. Epistemological Realism and the Basis of Scepticism*, Oxford, Blackwell.
Williams, M. 2001: *Problems of Knowledge*, Oxford, Oxford University Press.
Williamson, T. 2000: *Knowledge and Its Limits*, Oxford, Oxford University Press.
Wittgenstein, L. 1969: *On Certainty.* Ed. by G.E.M. Anscombe and G.H. von Wright, translated by G.E.M. Anscombe and D. Paul, Oxford, Blackwell.
Wright, C. 2004: "Warrant for Nothing (and Foundations for Free)?", *Aristotelian Society SupplementaryVolume*, 78, pp. 167-212.