When Science Confronts Philosophy: Three Case Studies

Eric Dietrich

ONLINE FIRST

Springer

Springer

ORIGINAL PAPER

When Science Confronts Philosophy: Three Case Studies

Eric Dietrich[1]

Abstract
This paper examines three cases of the clash between science and philosophy: Zeno's paradoxes, the Frame Problem, and a recent attempt to experimentally refute skepticism. In all three cases, the relevant science claims to have resolved the purported problem. The sciences, construing the term broadly, are mathematics, artificial intelligence, and psychology. The goal of this paper is to show that none of the three scientific solutions work. The three philosophical problems remain as vibrant as ever in the face of robust scientific attempts to dispel them. The paper concludes by examining some consequences of this persistence.

Keywords Computation · Frame Problem · Geometric series · Infinity · Knowledge · Relevance · Semantics · Skepticism · Zeno's Paradoxes

1 Introduction

This paper provides arguments that three famous philosophy problems are not solved or even dented by the best efforts of science and mathematics. The three problems are Zeno's Paradoxes, the Frame Problem, and Skepticism. The first and third are ancient. The Frame Problem in much younger; it is just now only a little over fifty years old. The conclusions will be the same for all three. The sciences (broadly construed to include mathematics) that allegedly solve each problem are all susceptible to the same objection: the empirical and theoretical methods used assume the existence of precisely what the philosopher calls into question. So, in each case, the scientist *begs the question* against the philosopher.

✉ Eric Dietrich
   dietrich@binghamton.edu

[1]   Philosophy Department, Binghamton University, Binghamton, NY, USA

2 Zeno's Paradoxes

Parmenides held that reality—the *What Is*—is a partless, continuous, unchangeable, immovable, unity that will exist forever, could never come into existence from "nothingness," and could never go out of existence. There is, finally, only pure crystalline being, now and forever. This is Parmenides's primary truth. A second truth (or path) concerns the *What Is Not*, which necessarily is not (i.e., it doesn't and cannot exist). Together, these two truths make up the *Way of Truth*. (In modern parlance we would probably call the Way of Truth, the *Way of* A priori, *Necessary Truth*.) Contrasting with the Way of Truth is the *Way of Opinion*. This way is merely the way of appearances, which merely *seem* to exist. This way is untrustworthy because it is often false or varied (one person has one opinion, another has the opposite opinion). Accordingly, it is the way of mortals, of contingent beings. This way is ugly and should be avoided, but, alas, it cannot, Parmenides held, due to its pervasiveness. At best, one can adopt a sort of detestable pragmatism toward the Way of Opinion. It may be that Parmenides held that the Way of Opinion is a pervasive illusion.[1]

Parmenides's philosophy is a shocking metaphysics and epistemology, and unsurprisingly it was not widely embraced. But one consequence of Parmenides's philosophy is still with us to this day, and continues to bedevil our understandings of motion, the relation between the many and the one, and the relation between thought and the world.

2.1 There is No Motion

Parmenides's metaphysics implies that there is no motion—nothing moves—another strange idea. But Parmenides had a student who was unusually bright: Zeno of Elea.[2] Zeno defended his teacher's metaphysics with at least four now-infamous motion paradoxes. On the power of these paradoxes, Bertrand Russell says:

> In this capricious world, nothing is more capricious than posthumous fame. One of the most notable victims of posterity's lack of judgment is the Eleatic Zeno. Having invented four arguments, all immeasurably subtle and profound, the grossness of subsequent philosophers pronounced him to be a mere ingenious juggler, and his arguments to be one and all sophisms (Russell and Bertrand 1903).

While it would be instructive to examine all four of Zeno's paradoxes, here we will only examine one, The Dichotomy. This paradox is easily accessible and widely

---

[1] See Palmer (2012) and Robinson (1968, pp. 108–109). The "Way of Truth" and "Way of Opinion" are often treated as their standard names, but other scholars have other names to the same effect. For example, Palmer calls the Ways of Truth and Opinion the Way of Conviction and the Way of Mortals, respectively (Palmer 2012).

[2] For the general ideas here in Sect. 2, I am indebted to Papa-Grimaldi's (1996) paper. I also draw on Benardete's (1964) conclusions.

known, and in attacks on Zeno's conclusion, the fallacy of begging the question against Zeno is easily seen.

Before we begin with the paradox, however, it will be worthwhile to clear up a potential source of confusion. Zeno didn't deny that we seem to experience motion, either by seeing something move or by seeming to experience our own movement. His argument was that while denying motion (and, more strongly, asserting Parmenides's unitary *What is* was the correct metaphysics) led to conflicts with perception (we do seem to see plenty of motion and we do seem to see a multitude of things), asserting that motion existed (and, more strongly, that Parmenides was wrong) led to conflicts with *logic*—which is far worse.[3] After all, perception is well-known to be fallible. Zeno's position, then, is that motion belongs to the ugly *Way of Opinion*. Finally, it is worth stressing that all of Zeno's paradoxes are clashes between what we seem to experience and our logical reasoning. Zeno naturally insisted that the latter takes precedence over the former, 'naturally', because he didn't want to violate the law of non-contradiction.

The Dichotomy Let's have Achilles run from a stationary point *A* to another stationary point *B*, a distance, let's say, of a 100 meters. Zeno asserts that Achilles cannot run this distance. There are two versions of The Dichotomy: the progressive version and the regressive version. Here, we are only interested in the progressive version. In the progressive version, Achilles cannot get to *B* (in the regressive version, he cannot even get started; he can't leave *A*). In order to make it to *B*, Achilles has to first go halfway to *B*: 50 meters. 50 meters remain. Now to get to *B*, Achilles must again go halfway to *B*: 25 meters. 25 meters remain. Now to get to *B*, Achilles must go yet again halfway to *B*: 12.5 meters. 12.5 meters remain. And so forth, ad infinitum. We see, then, that Achilles has to cover an infinite number of decreasing halfway distances in order to get to *B*. So, there is always some distance to cover before he can get to *B*. So, he cannot get to *B*. (Though this argument seems to break up Achilles's run to *B* into discrete distances, this is an artifact of the presentation of the argument. Achilles's attempted run to *B* is smooth and continuous… and never ending.)

2.2 Does Math Solve Zeno's Paradoxes?

It is widely held that the sum of a specific geometric series dispels the progressive form of the Dichotomy (e.g., Salmon 1975). Euclid knew well this sort of series and others like it (and so did others before Euclid). Euclid derived a general formula for the expression of the sum of any geometric series (discussed below). If the sum converges on a single number, the series is said to *converge*. If the series did not sum to a specific number, but grew beyond all bounds, it is said to *diverge*. A much clearer and cleaner picture of the convergence or divergence of a geometric series only emerged after mathematicians Cauchy, Dedekind, Weierstrass, and others finally placed the foundations of the calculus on a firm conceptual footing,

---

[3] See for example, Robinson, 1968, p. 128.

repairing Newton's and Leibniz's initial foray into the infinitesimal. The relevant notions Cauchy and company gave us include *limit*, *approaching a limit*, *converging to a limit*, and *approaching infinity*, as well as others.

The series we are interested in is the geometric one below, which we will dub the *Z Summation*. The crucial, provable property of this sum is that it is an infinite summation that *converges* to a finite sum.

The Z Summation

$$\frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \frac{1}{16} \cdots$$

Here, we add up all the numbers between 0 and 1 of the form $\frac{1}{2^n}$, where n takes on all numbers 1, 2, 3, 4, 5,…. Intuitively, this infinite summation sums to 1: the sum will never grow beyond 1 because what's always added is one-half of the remaining distance to 1, and the sum clearly gets closer and closer to 1.

Of course, we don't actually do an infinite number of additions—we can't. Rather, the math allows us to *derive* the sum of the Z Summation, namely 1, by considering the sum in the limit, as n approaches infinity. The details are crucial for both the alleged refutation of Zeno's conclusion and for Zeno's defense of his conclusion, so here, briefly, is how this works.[4]

To begin, here is the formula for the sum of the first n terms of a series:

$$(1) \quad S_n = \frac{a(1 - r^n)}{1 - r} \tag{1}$$

S_n is our sum of the first n terms of any series, a our starting number, and r our constant ratio used to control the growth of the series. Deriving this formula took insight and cleverness. (Apparently, it is not known who first derived this formula—someone who lived well before Euclid; Euclid's proof that this formula will work is different from the one used today in most calculus books (Kaplan and Kaplan 2003, p. 90). From (1), it is easy to derive (2), the sum of an infinite series, provided that $|r| < 1$, for then as $n \to \infty$, $r^n \to 0$.[5] We thus get:

$$(2) \quad \lim_{n \to \infty} S_n = \frac{a}{1 - r} \tag{2}$$

This is read: "the limit of S_n as n goes to infinity is $a/(1 - r)$." Again, S_n is our sum, a our starting number (in the Zeno case, 1/2, since Achilles runs half of each of the remaining distances), and r, also here 1/2, controls the growth of the series. The series converges (produces a sum) because the absolute value of r is less than 1: $\left|\frac{1}{2}\right| < 1$.

(2) then is how we get the finite sum of the infinite Z Summation: making the relevant substitutions gives us $\frac{\frac{1}{2}}{\frac{1}{2}}$, which equals 1, our desired result.

---

[4] The proofs for what is to follow can be found in any calculus or real analysis book, and online in various places, such as *Wikipedia*: https://en.wikipedia.org/wiki/Series_(mathematics).

[5] If $|r| \geq 1$ then (1) diverges, i.e., goes to infinity.

Now, the argument against Zeno is this. The Z summation looks a lot like the progressive form of the Dichotomy (as many have noted, again see Salmon 1975). If this summation equals 1, then by covering an infinite number of decreasing halfways (as described above), Achilles can indeed race to the end of one race course. Paradox dispelled. The alleged power of this mathematical dispelling can be summarized in the following exchange between Zeno and the Fans-of-Using-Math-to-Dispel-Zeno's-Paradoxes:

Zeno:

1. Assume Achilles attempts to run from point *A* to point *B*.
2. Then, by logical reasoning, we see that he can never complete his run (because he has to cover an infinite number of half-distances).
3. But he does *seem* to complete the run.
4. This results in a contradiction between the Way of Opinion and the Way of Truth.
5. Clearly the Way of Truth is the Way to follow, for it is the more certain.
6. Therefore, Achilles cannot complete the run.

Fans of Math:

1. Assume Achilles is running from point *A* to point *B*.
2. This results in no contradiction between experience (Way of Opinion) and reasoning (Way of Truth) because the Z Summation is finite—1, to be exact. We know this because the Z Summation converges, as shown in (2).
3. There is now no roadblock to letting the compelling nature of our experience of motion guide us to acknowledging motion's real existence.
4. Therefore, Achilles really is moving, and motion is perfectly okay.

Zeno is unmoved. He gives the following argument in reply to the Fans of Math:

The derivation of the sum of the Z Summation confirms my conviction, not yours. As you Math Fans freely admit, this infinite summation cannot be carried out directly in the physical world simply because it requires an infinite number of 1/2's to be actually summed, which is impossible. We know that the Z Summation converges to 1 not because we do the actual summing, but because *in our minds,* we make the appropriate *inferential leap* from a finite amount of summing in (1) to the full infinite summation using (2). The brilliance of (1) and (2) is that we can mentally derive the sum of the Z summation (or any other appropriate geometric series) without actually doing the impossible, summing up the infinite additions. Indeed, this is the point of (1) and (2). But Achilles, in running from *A* to *B* can't "mentally derive" or "inferentially leap" to *B*—he is running, not thinking. I do not doubt that the Z summation sums to 1. I merely deny that this sum can be reached by doing the actual summing of all the 1/2's. This sum can only be reached by doing mathematics. And you actually agree with me on this. Again, Achilles is not doing math, he's running (or appears to be). You are merely assuming that Achilles is somehow doing what your mathematics shows to be true. By mak-

ing this assumption, you are further assuming that Achilles is moving. This latter assumption, and so the former one, beg the question against me. You are assuming precisely what I am denying. It is easy to win an argument if you assume what your opponent is denying. I conclude that the sum of the Z Summation (which is not in dispute) is *irrelevant* to my arguments. My paradoxes remain even though we agree on the sum of the Z Summation and agree that the Z Summation looks a lot like the progressive form of the Dichotomy.[6]

This reply to the Fans of Math seems decisive. The Fans either have to agree with Zeno or beg the question against him, a vacuous victory. Zeno's paradoxes prevail.

2.3 In Sum

We have here an example of the pattern of interest. This confrontation between science (mathematics) and philosophy leaves the philosophical point unfazed—indeed, untouched. Mathematics confronts Zeno's paradoxes, changes the rules, assumes what Zeno explicitly denies, and then proclaims that it has refuted the paradoxes. Math has done nothing of the sort. The sum, 1, can be mentally, mathematically inferred, but one cannot run to it, count to it, nor traverse a space to it. Achilles is an actual physical being. No actual physical being can do an actually infinite task—which is why we need (1) and (2). Yes, somehow Achilles appears to run from point *A* to point *B*. But that's just the perversity of the Way of Opinion…. We were warned to stay clear of it.

3 The Frame Problem

Humans are probably the smartest animals on planet Earth. The Frame Problem begins here: What enables us to be so smart? Aristotle intuited, dimly, what is perhaps the key: "The animals other than [humans] live by appearances and memories, and have but little of *connected experience*; but the human race lives also by art and reasonings" (Aristotle, Book 1, Part 1, *Metaphysics*, emphasis added; Aristotle intends to include only rational reasonings—for him there was no other kind of reasoning). The Frame Problem is the problem of how we so robustly connect our experiences. This no doubt is a bit cryptic; the Frame Problem takes quite a bit of set-up to explain.

3.1 Defining the Frame Problem

The Frame Problem is ultimately about *relevance*—semantic, cognitive relevance. The problem of what is relevant to what runs deeply throughout large parts of

---

[6] The foundations of the argument presented here can be found in: Papa-Grimaldi's "Why mathematical solutions of Zeno's paradoxes miss the point" (1996) and Benardete's *Infinity*, especially, p. 12 (1964), where Benardete argues that Zeno's paradoxes remain unsolved and unscathed to this day.

philosophy of language, philosophy of mind, metaphysics, and epistemology. However, the Problem's central two *technical* issues involve *change* and *updating*. All living things live in changing environments. Among thinking things, some of the most important changes that they must deal with are changes they themselves bring about via their actions. However, more than change is involved. To successfully handle any change, thinking things must update what they believe and expect after experiencing the change. The major difficulty wrought by the Frame Problem is questioning how, and how far, this updating should be carried out. This is best explained by an example.

If you put a pot on the stove to boil water, you need be aware that the lid, including its handle, is likely to get hot, too; so you can't just grab the lid's handle with your fingers to check on the boiling water or to put in, say, some rice or pasta. Putting a pot on the stove to boil water is a change, and this change requires updating one's belief about the temperature of the lid and its handle (note, however, the temperature of the pot itself quite naturally gets updated in your mind—you would never just grab the pot between your two hands). Failure to so update the temperature of the lid will result in burned fingers. This can be put in terms of *relevance*: the temperature of the pot is relevant to the temperature of the lid, so if the former changes, so does the latter.

But now there are also the problems concerning what you can afford *not* to update if considered, and what you can simply *ignore* utterly. Consider two further cases.

Case one: With a pot on the stove, your belief that your car has four tires should not be updated for the simple reason that pots of boiling water on your stove do not affect how many tires your car has—the two topics are not relevant to each other. (Note: one can make up a story where the two *are* related, but that story will have a very low probability of occurring).

Case two: Given the pot of boiling water on the stove, should you update your belief that Antarctica is still cold? It depends. On the one hand, arguably Yes, since your use of your stove and the use of stoves all around the world are helping raise the global temperature by burning fossil fuels, which, in turn, is warming Antarctica. But on the other hand, No, since Antarctica is still cold. Perhaps it is less cold by a tiny fraction of a degree because of your boiling water, but this fraction is probably not even measurable.

Now here is the key. Until they were mentioned, you almost certainly would not have even considered updating your beliefs about your car's tires, nor the ones about Antarctica's temperature. In fact, you would not have considered the vast majority of your beliefs for updating; nor should you have, because you have an enormous number of beliefs; considering them all would take a lot of time. But hidden in your enormous number of beliefs are some that are in fact *relevant* to your pot of boiling water, but which, up to now, you would have considered *irrelevant*.

Aristotle's claim above was that part of humans' robust intelligence is how we connect our experiences. Connection is captured by the notion of relevance. Importantly, which of your beliefs are relevant to any current change you are experiencing is not fixed. Relevance itself changes as the world changes. Beliefs once thought to

be irrelevant become relevant, and vice versa. So, your tire belief should not even be considered when you update your beliefs given your pot of boiling water. But your Antarctica belief is a candidate for such updating (at least arguably).

And now we have the Frame Problem. To find the beliefs relevant to a given change, you do, apparently, have to canvass your *entire* belief store—a vast collection of beliefs.[7] Some will be more obvious than others, and some more relevant than others, but the entire store has to be canvassed. So, there you stand examining and examining while the water in your pot boils away.

3.2 Victimized by the Frame Problem: Heuristic Updating

Something has gone badly wrong. Given any change, we do not and cannot examine the whole of our belief store because it is too big. But this is precisely why we fall victim to the Frame Problem. Losing your keys, wallet, or cell phone is often an instance of the Frame Problem. You put your keys down, but not in the spot you normally put them; you don't update your beliefs correctly, and *oops!* lost keys.[8]

Being a victim of the Frame Problem is just part of living. If your mistake is bad enough, the consequences will be, too. But usually we find our keys. The fact that we are often victims of the Frame Problem strongly makes the case that given an experienced change, what humans and other animals do is use the following heuristic for updating: *update those beliefs that seem now to be obviously relevant, while ignoring the rest.*

3.3 The Official Frame Problem

To complicate matters, the above description of the Frame Problem is not the official version, according to AI scientists. This takes some explaining. The Frame Problem is a complicated problem with a tortured history. And it is not possible to define it succinctly and easily. In fact, this was a big part of what the struggle over the Frame Problem was about.

The Frame Problem began life in an infamous 1969 paper by John McCarthy and Patrick Hayes ("Some Philosophical Problems From the Standpoint of Artificial Intelligence"). In this paper, McCarthy and Hayes, both computer scientists (this matters), introduced the term "Frame Problem" to denote a seemingly narrow

---

[7] It does not matter for our purposes whether most of one's beliefs are made explicit and stored as such or whether the beliefs are left implicit, to be derived from a relatively small set of more active beliefs. Either way, an enormous number of your beliefs must be canvassed and analyzed for relevance to the current change, *for each change*. It is also worth stating that it is here assumed that one's belief store, no matter how large it is, is finite.

[8] Global warming is a stunning example of the Frame Problem on world-wide scale. Few updated their beliefs and understandings about what life would be like if billions of cars were driven around every day pumping, together with factories of all kinds, large quantities of CO_2 into the air while rain forests were razed around the globe. True, in the 1960s some concerned scientists finally warned the world that global warming was a likely result of industrialization, but few in government or industry listened, and anyway, by then it was too late. The evolution of Antibiotic-resistant pathogens is another example.

logic problem that arose while the two worked to develop a logic for modeling reasoning about change.[9] But just nine years later, in 1978, the Frame Problem had become "an abstract *epistemological* problem," (see Dennett 1978, p. 125, emphasis in original). Then, in 1987, the philosopher Jerry Fodor equated the Frame Problem with "the problem of how the cognitive mind works" (Fodor 1987, p. 148). He then claimed that understanding how the mind works requires unraveling the nature of inductive relevance and rationality (ibid.). So, Fodor was saying that solving the Frame Problem would be figuring out how inductive relevance and rationality operate. From here, the Frame Problem continued to grow and expand until it covered vast areas of philosophical research. The Frame problem was revealed to be a serious and deep philosophical problem, and therefore, probably completely intractable; hence humankind's heuristic solution to it discussed at the end of Sect. 3.2.[10]

In fact, Fodor thought the Frame Problem was so serious that it explained why AI had failed thus far and was going to continue to fail. Fodor said:

> We can do science perfectly well without having a formal theory of [which ideas or events are relevant to each other]; which is to say that we can do science perfectly well without solving the Frame problem. That's because doing science doesn't require having *mechanical* scientists; we have *us* instead. But we can't do AI perfectly well without having mechanical intelligence; doing AI perfectly well just *is* having mechanical intelligence. So we can't do AI without solving the Frame Problem. But *we don't know how* to solve the Frame Problem. That in a nutshell, is why, though science works, AI doesn't. (Fodor 1987, p. 148; emphases in original).

But of course, AI researchers and their philosophical allies disagreed with Fodor. Hayes (1987) gives a definition of the Frame Problem. He asks us to consider a case where someone goes through a door from room 1 to room 2. Hayes says that we want to be able to prove (in some logic) that when an agent (a thinking thing) goes from room 1 to room 2, then the agent is *in* room 2. (Again, this might seem obvious to human readers, but to be obvious to a computer, the computer has to be programmed correctly.) To get this conclusion that going from room 1 to room 2 puts one *in* room 2, we need *an axiom* to this effect. No problem. We simply add it in. Hayes then says:

---

[9] Why is it called the *Frame* Problem? McCarthy and Hayes say: "In the last section of part 3, in proving that one person could get into conversation [over the phone] with another, we were obliged to add the hypothesis that if a person has a telephone he still has it after looking up a number in the telephone book. If we had a number of actions to be performed in sequence we would have quite a number of conditions to write down that certain actions do not change the values of certain fluents. In fact with *n* actions and *m* fluents we might have to write down *mn* such conditions. [A fluent is a sentence or predicate in the logic being used that is a condition, property, or state of affairs that can change over time.]" Then McCarthy and Hayes say: "We see two ways out of this difficulty. The first is to introduce the notion of [a] frame.… A number of fluents are declared as attached to the frame and the effect of an action is described by telling which fluents are changed, all others being presumed unchanged.".

[10] The Frame Problem is intractable because it is a philosophy problem. I say this because I hold that philosophy makes no progress and cannot solve its major problems. See Dietrich (2011).

But here at last is the frame problem. With axioms [like the one we are adding about changing rooms], it is possible to infer what the immediate consequences of actions are. But what about the immediate non-consequences? When I go through a door, my position changes. But the color of my hair, and the positions of the cars in the streets, and the place my granny is sitting, don't change. In fact, most of the world carries on in just the same way that it did before… But since many of these things CAN change, they are described in our vocabulary as being relative to the time-instant [of when the room change occurred], so we cannot directly infer, as a matter of logic, that they [remain] true [after I change rooms] just because they were [true before I changed rooms]: This needs to be stated somehow in our axioms (p. 125).

Then Hayes, continuing, points out:

In this ontology, whenever something MIGHT change from one moment to another, we have to find some way of stating that it DOESN'T change whenever ANYTHING changes. And this seems silly, since almost all changes in fact change very little of the world. One feels that there should be some economical and principled way of succinctly saying what changes an action makes, without having to explicitly list all the things it doesn't change as well; yet there doesn't seem to be another way to do it. *That* is the frame problem (p. 125, emphasis in original).

What Hayes wanted, what all Frame Problem AI researchers wanted, was some epistemic and metaphysical *stability principle* or *principle of epistemic inertia* which worked *in all cases and for all tasks no matter how complicated.* But philosophy was exactly pointing out that such a principle was not possible because we cannot know ahead of time all of what is relevant to a given change or action. For starters, changes change relevance relations. So, we actually have to make the change and then wait and see what other things change as a result. If it were otherwise, we could do all of science from our armchair armed only with pencil and paper.

3.4 The "Two Narratives" Interpretation

Hayes disagreed with Fodor, Dennett, and other philosophers that the Frame Problem was profound and that it was deeply connected with how the mind works and knows whatever it knows. Hayes replied to Fodor:

The term "frame problem" is due to John McCarthy, and was introduced in McCarthy and Hayes (1969). It is generally used within the AI field in something close to its original meaning. Others, however, *especially philosophers*, sometimes interpret the term in different ways… In this short paper I will try to state clearly and informally what the frame problem is, distinguish it from other problems with which it is often confused, briefly survey the currently available partial solutions, and respond to some of the *sillier misunderstandings* (1987, p. 123; emphasis added; as Hayes's paper proceeds, the language gets sharper).

By sillier misunderstandings Hayes means, of course, *philosophical misunderstandings*.

The eventual compromise was to say there were *two narratives* about the Frame Problem. The philosophy narrative says that the Problem is as deep as metaphysics and the nature of rationality, and as complex. Therefore, understanding how humans heuristically avoid the Frame Problem some significant amount of time is probably one of the keys to understanding the human mind (see Fields 2013). Opposing this narrative, we have the AI narrative, which says the Problem is only a challenging technical problem in the logic of reasoning about change. And, if that weren't different enough, on this latter narrative, the Frame Problem has been *more or less solved*! Shanahan (2016), who adopts completely the two-narrative narrative says:

> … [A] number of solutions to the technical frame problem [the logic version] now exist that are adequate for logic-based AI research. Although improvements and extensions continue to be found, it is fair to say that the dust has settled, and that the frame problem, in its technical guise, is *more-or-less* solved (emphasis added).

Philosophers impressed with the Problem will point out that "more or less solved" means *not solved.* All the logic solutions to the Frame Problem invoke some sort of strong restrictions or circumscriptions to the target logic so that accounting for unintended effects of changes due to actions is kept manageable. There is no general algorithmic or logical solution to the AI interpretation of the Frame Problem; nor is there any other kind of solution. There are only heuristic avoidances of the Problem that work locally, and only more or less well; see again the heuristic updating "solution" at the end of Sect. 3.2. So, the AI interpretation of the Frame Problem is not so much a solution to the problem as a finely crafted way to avoid the problem. That just leaves the philosophical interpretation.

3.5 The Profundity of the Frame Problem

A particularly well-developed argument for the above conclusion is found in Fields' 2013 paper. Fields concludes, upon considering the relevant neuroscience, that the Frame Problem and the problem of object re-identification (How do we re-identify objects as persisting things from one context to another?) are *equivalent*. Object re-identification is widely thought to be "solved" in part by pre-motor fictive (made up) causal histories, clearly a heuristic process that can only work some of the time. Almost all of AI and large areas of cognitive science (notably the study of cognition) assume that objective re-identification is unproblematic. Clearly this assumption will have to be abandoned if the relevant disciplines are to have any hope of unraveling how humans successfully avoid the philosophical version of Frame Problem—when they do. Finally, Fields points out that object re-identification is fundamentally *analogical* in nature. Hence this analogical character will have to be deployed both in AI approaches to the Frame Problem as well as to understanding how humans deal with the Problem.

But of course, analogy-making in the mind is a creative process that alters the very mental representations one is using to think about those changes one is causing in one's environment and the objects that need to be re-identified (see Dietrich 2000, 2010, and also see Connell and Lynott 2014). What guarantees that those internal representations still represent what their immediate precursors represented? Further use of analogical thinking leads to an infinite regress. Hence the Frame Problem, as well as object re-identification, emerges again, inside the minds of those trying to avoid it in their daily lives. It is clear that the mind itself begs the question when it comes to the Frame Problem: It simply assumes that certain representations re-identify objects. If it didn't, an infinite regress would result, and we'd all come to screeching halt.

We now have here the same pattern we had with Zeno's Paradoxes. Those seeking to avoid the philosophical gravity of a certain problem simply deny that such gravity exists and declare that the problem commits a technical error correctable with enough fancy logical or mathematical machinery. The philosophers reply that deploying this machinery is entirely question-begging because the machinery derives from asserting the very claim they (the philosophers) deny. It is easy to more or less solve the Frame Problem if one denies its profundity and ubiquity. That the Problem returns down deep in the mind is hardly surprising to philosophers who take it and philosophy seriously.

The Frame Problem, therefore, is here to stay. Humans deal with it as best they can, just as they do when re-identifying objects in their perceptual environments. Understanding fully, not more or less, the persistence and gravity of the Frame Problem requires understanding deeply both the mind and the fundamental nature of reality, philosophical problems both.

4 Experimentally Refuting Skepticism[11]

Skepticism, the view that we know nothing, is both outlandish and persistent. Instead of attacking skepticism directly, a new strategy would be to use human emotional and epistemic preferences to explain why such an implausible view has such staying power. This strategy would then help *explain away* skepticism, rather than attempting to refute it directly. The claim would then be that we humans have a natural but unjustified bias against claiming knowledge in the very cases skepticism depends on. This is Turri's strategy in his *Skeptical Appeal: The Source Content Bias* (2015).

The classical skeptical argument aims to prove that insofar as we do not know that a skeptical hypothesis does not hold, our ordinary beliefs about the world fail to constitute knowledge. The argument is typically cast in the following form. First, let *O* be a proposition one would ordinarily take oneself to know (e.g. that the external world exists or that one has hands) and let *T* be a tricky, skeptical hypothesis such as that one is dreaming or being systematically deceived. *Not-T*, accordingly, is the denial of the skeptical hypothesis. The form of the skeptical argument is then:

---

[11] Thanks to Anna Tsvetkov for co-authoring parts of this section using her surperlative 2018 Department of Philosophy Honor Thesis (Binghamton University) entitled "Saving Skepticism: A Response to Turri."

If I know that O, then I know that not-T.
But I do not know that not-T.
Therefore, I do not know that O (Turri 2015, p. 309).

For example, a skeptic might raise the brain-in-a-vat hypothesis (*T*): one might be a disembodied brain in a vat that receives electrical stimulation from a powerful computer. The stimulation induces a set of false experiences of an external world qualitatively indistinguishable from the experiences one enjoys now. Suppose Jones ordinarily knows that she has a physical body (*O*). The skeptic argues that if Jones knows this, then Jones knows that the brain-in-a-vat hypothesis is wrong (*not-T*). The skeptic then points out that Jones does *not* know that she is not a brain in a vat; she cannot rule out the brain-in-a-vat hypothesis because Jones's experiences remain the same regardless of whether or not the skeptical hypothesis is realized. From these premises Jones is led to accept the conclusion: "I do not know that I have a physical body." This argument works for any ordinary knowledge claim. The skeptical argument utilizes intuitively plausible premises to lead us to profound and startling conclusions: no knowledge claim is secure.

Many strategies have been deployed against the skeptical argument. Nozick (1981) and Dretske (1970, 1971), reject epistemic closure and consequently, the first premise of the skeptical argument. Other philosophers reject the second premise of the skeptical argument in favor of an appeal to common sense or to semantic externalism (see Moore 1962; Putnam 1975, respectively). Contextualists, such as Lewis (1996) and DeRose (2009), argue that the skeptic's argument garners putative force by simply raising the semantic standards for knowledge. This last strategy of attacking the force of the skeptical argument is a version of the strategy Turri uses.

After conducting a series of experiments (2015), Turri concludes that the potency of the skeptical argument can be explained away as a byproduct of human psychology. He maintains that humans possess a *source-content bias*: we are biased against classifying negative inferential beliefs as knowledge, where a negative inferential belief is a derived belief that something is not the case.[12] The second premise of the skeptical argument—a denial of knowledge of a negative inferential belief—achieves acceptance by preying upon our evaluative bias. Turri concludes that skepticism is an illusion that results from our psychology.

---

[12] It is important to note that the negative inferential belief is *no*t not-T. Turri doesn't elucidate this. The belief is implicit in the first premise of the skeptical argument, derived here, where K is a modal knowledge operator meaning "It is known that":

(1) $O \rightarrow$ not-T
(2) $K(O \rightarrow$ not-T$)$
(3) $\therefore K(O) \rightarrow K($not-T$)$

The negative inferential belief is explicit in the conclusion here (which is the first premise in the skeptical argument). Turri's experimental results seem to show that we are reluctant to classify the inference from O to not-T as knowledge (the second premise), and so we are reluctant to say the consequent of the conclusion is true. Hence we are led, according to Turri, to deny that we know that not-T. (The conclusion follows from the second premise via the Distribution Axiom of the modal logic K (Kripke): $\Box(p \rightarrow q) \rightarrow (\Box p \rightarrow \Box q)$.)

4.1 Turri's Experiments: Implicating a Source-Content Bias

Turri is particularly interested in locating the *appeal* for skeptical arguments outside the power of skepticism and the relevant logic itself. He maintains that no one reasonably accepts skeptical arguments within one's daily life. One does not go about contending or even worrying that one does not know that other people have minds, that one is not dreaming, or that one is interacting with an external world. Yet skepticism endures. Skeptical arguments appeal to us when they are considered, at least in certain settings, like in an epistemology class. Turri examines whether the cause of the appeal and force of skepticism can be attributed to preferences in human psychology. Namely, do humans possess a bias which renders forceful an otherwise impotent, skeptical argument?

Turri hypothesizes and claims to experimentally show that humans possess a *source-content bias*. The bias is two-fold. First, humans evaluate inferential beliefs more harshly than perceptual beliefs. People more readily classify beliefs derived from direct perception as knowledge than beliefs arrived at by inference. The *source* of a belief, therefore, affects whether or not one classifies one's belief as knowledge. Second, humans evaluate negative inferential beliefs more harshly than positive inferential beliefs. A belief stated in a positive form—that something *is* the case—is more readily accepted as knowledge than a belief that states that something is *not* the case. Hence, the *content* of a belief affects rates of knowledge ascription and denial. Turri hypothesizes that the second premise of any instantiation of the skeptical argument form (at the beginning of Sect. 4) is appealing because of an interaction of these two effects: the inferential source and negative content of the relevant belief.

To test his theory, Turri conducted two experiments. Turri concludes: "… the skeptic has simply alighted on a class of beliefs that we are antecedently inclined to evaluate especially harshly" (2015, p. 316); and finally, "It is not… some deep fact about the nature of knowledge or the skeptic's ingenuity that invites skepticism. It is our psychology" (2015, p. 320). We humans only find the skeptical argument to be forceful, and accordingly, worry about skepticism, when we are presented with the skeptical argument and are misled by our bias. Otherwise, as Turri claims, we easily renounce skepticism.

For convenience in the following discussion, let us recast Turri's argument as follows.

1. People evaluate inferential beliefs more harshly than perceptual beliefs.
      [Per Turri's experimental findings.]
2. People evaluate an inferential belief more harshly when its content is negative.
      [Per Turri's experimental findings.]
3. If (1) and (2) then the source-content bias exists.
4. Therefore, the source-content bias exists.
5. The second premise of the skeptical argument is a harsh evaluation of a negative inferential belief.

6. If the source-content bias exists and the second premise of the skeptical argument is a harsh evaluation of a negative inferential belief, then the appeal and force of the skeptical argument is a byproduct of our psychology.

   (Skeptical arguments simply prey upon our evaluative bias against classifying negative inferential beliefs as knowledge. According to Turri, the bias can also "explain the classical skeptical argument's force" [2015, p. 310].)

7. Hence, the appeal and force of the skeptical argument is merely a byproduct of our psychology.

8. If the appeal and force of the skeptical argument is a byproduct of our psychology, then skepticism only results from our psychology.

9. Therefore, skepticism only results from our psychology.

10. Therefore, skepticism is not tenable.

There are now three objections to Turri's project.

4.2 Objection 1: Experiments Beg the Question Against the Skeptic

Has skepticism been explained away? A skeptic is going to be unimpressed with Turri's experiment-based argument because the argument begs the question against the skeptic.

Consider the skeptical argument again. It is important to note that the argument generalizes. That is, the argument can be employed not only to show that one does not know that an ordinary belief one holds is true, but that *none of the* beliefs one holds can be known to be true. But Turri must begin with the assumption that something *is known*—namely, scientific knowledge—to argue against skepticism.

Turri's paper follows in the tradition of naturalized epistemology. According to this view, philosophical questions concerning the nature of knowledge, evidence, justification and so forth, can be empirically tested and answered. Epistemology is cast within the experimental domain, alongside skepticism. The success of Turri's approach hinges on the possibility for skepticism to be addressed *within* the field of psychology. But skepticism casts *all* knowledge under doubt, including scientific knowledge. To address skepticism by accepting and utilizing the putative scientific knowledge in one's repertoire is to beg the question and egregiously so.[13] How can one know one's scientific knowledge is true? More generally, how can one know that the scientific enterprise produces knowledge? To address these concerns with further empirical evidence is to give a circular argument: science produces knowledge

---

[13] Consider the countervailing view. Quine, a proponent of naturalized epistemology, argues: "skeptical doubts *are* scientific doubts" (1975, p. 65, emphasis added). That is to say, skeptical doubts are akin to an awareness of illusions. Illusions can only arise against the backdrop of certain beliefs formed by science. Because of their empirical origins, one can employ scientific methods to neutralize skeptical concerns without begging the question. We reject this argument because of its reliance on Quinean empiricism and because of its circular nature. Skepticism calls our fundamental beliefs into questions. One cannot justify one's basic beliefs by the very belief-forming methods called into question by raising the skepticism issue. See Fumerton (1994) for a similar objection.

because scientific results (i.e. evidence) constitute knowledge. But each piece of purported evidence is laden with skeptical doubt.

Turri would reject this skeptical reply to his argument. First, he assumes that empirical investigations can, indeed, provide us with conclusive evidence to constitute knowledge. Although he remarks in his paper that he intends to avoid making, "sweeping conclusions on the nature of knowledge", elsewhere he writes, "it is not, in the first place, some deep fact about the nature of knowledge or the skeptic's ingenuity that invites skepticism. It is our psychology" (2015, p. 320). But to state that it is not the nature of knowledge that invites skepticism, but rather a psychological bias *is* to make a sweeping conclusion on the nature of knowledge. Recall that Turri suggests that, without the bias, one could know that the skeptic's alternative hypotheses are not true. In other words, skepticism only results from an evaluative bias. This is a knowledge claim derived from experimental results, which the skeptic will deny.

Second, Turri presupposes *realism*, the view that there is a mind-independent external world. In order for Turri to conduct empirical investigation he must make several epistemological and metaphysical commitments. He must accept the existence of an external world and the validity of epistemic norms (e.g., the scientific method) that allow for findings that reflect a true (or approximately true) account of a natural world. But the skeptical hypotheses threaten realism and the possibility for *any* knowledge of an external world. To defeat skepticism by starting with presuppositions of an external world is again to beg the question against skepticism.

4.3 Turri Counters

Perhaps the best way for the experimental-anti-skeptic to respond to the objection of begging the question is to shift the burden of proof. That is, Turri might argue that he, himself, is not making any question-begging assumptions by employing scientific methods. If the skeptic saddles scientific methodology with doubt, it is she who thereby takes upon herself the burden of proof. That is to say, the skeptic must provide a *successful* argument for the acceptance of her skeptical conclusions. The anti-skeptic might argue no such argument is given. As a result, the question-begging objection against Turri ought to be rejected.

As it stands, this rebuttal is powerful. Recall the skeptical argument:

If I know that O, then I know that not-T.
But I do not know that not-T.
Therefore, I do not know that O (Turri 2015, p. 309, emphasis added).

The first premise is uncontroversial. But the skeptic must persuade the experimentalist of the second premise to reach the skeptical conclusion. That is, the anti-skeptic must concede that *she does not really know* that she is not a brain in a vat or that she is not being systematically deceived by an evil demon, and so forth for the skeptical argument to prove successful. But the anti-skeptic, naturally, does not find any such premises forceful. For she finds greater intuitive appeal in the doctrines of realism and the proposition *that there is an external reality* than in any claims to the

contrary. This, after all, is the thrust of such Moorean facts: they are propositions one knows better than the premises of any competing philosophical argument.[14] The skeptic has failed to provide any reasons for the anti-skeptic to accept her conclusion. Indeed we might make the stronger claim: no such reason can be given. The skeptic seems to have lost from the get-go.[15] She cannot, in principle, provide a successful argument to the anti-skeptic. No reasons will persuade the anti-skeptic of the pivotal second premise.

Attractive as this objection might be, there are two central motivations to resist it. The first motivation, derived from Pryor (2000) is as follows. Recall Turri's project takes the central aim of explaining away skepticism as the byproduct of a psychological bias. A bias which is inherent in all of us: skeptics and anti-skeptics, alike. Accordingly, Turri takes on what Pryor dubs to be an *ambitious anti-skeptical project*. That is to say, Turri's project aims to give a robust and universal answer to the problem of skepticism. Skepticism can be explained away as a psychological illusion, of sorts. Upon reading Turri's results, both the skeptic and the anti-skeptic ought to come to recognize the putative force of the skeptical argument merely stems from an illusion. Simply put: Turri intends to explain away skepticism *for the skeptic*.[16] But he cannot do so by employing methodologies and making stipulations the skeptic would not allow on pain of begging the question. Because Turri has taken up an anti-skeptical project of the ambitious kind, he comes to bear the burden of proof against the skeptic. Otherwise, his project is for naught.

The second motivation is simple. There is an important sense in which the skeptic cannot be said to beg the question against the dogmatist, or anti-skeptic. That is, the skeptic has the following response available to her: she can recommend suspension of all belief, either positive or negative. She can argue, like Sextus Empiricus, that she is neither asserting nor denying her position in deploying the skeptical argument; rather, she is in a state of *epoche* where she makes no epistemological commitments (see Sextus's *Outlines of Pyrrhonism*). Using this strategy, our skeptic claims that in *epoche* she seems to be speaking (making certain noises), and that if there are anti-skeptics in hearing range of her voice, the rational and logical force of her alleged claims is a problem for them, not her. Maintaining such *epoche* might be difficult for the skeptic, but that is due merely to habit, and not to any deep fact about knowledge or the world.

Now because the skeptic resides in *epoche* she has no positive epistemological commitments to refute. She cannot be said to beg the question. For she makes no assertions. The skeptic enjoys a tranquility of mind precisely because she suspends all belief. Indeed it is Turri who bears a burden of proof because he, unlike the skeptic, *is* asserting propositions and so bears the relevant epistemic and metaphysical commitments.

---

[14] Lewis (1996).

[15] See Kelly (2005) for a similar objection.

[16] For otherwise, Turri's paper bears an uninteresting thesis. Surely, there is little (perhaps no) motivation in shedding light on a tremendous bias towards accepting skeptical beliefs to those (the anti-skeptics) who do not accept such beliefs in the first place.

Though the begging-the-question charge against Turri is quite powerful, most will reject it (if only for pragmatic reasons) because at a minimum, it stifles engaging with the deeper parts of Turri's anti-skepticism project. It is to these, we know turn.

4.4 Objection 2: Two Kinds of Negative Inferential Beliefs

We need a distinction, one missed by Turri, because not all negative inferential beliefs concerned with knowing are equally susceptible to Turri's experiment-based argument. Consequently, not all negative inferential beliefs are equally powerful in giving us skepticism.

Distinguish between *quotidian* negative inferential beliefs (about knowing) and *world-changing* negative inferential beliefs (about knowing). Turri's experiments are exclusively about the first kind, negative inferences we draw in our ordinary lives. In one experimental set-up, Person-1 makes inferences about whether his car has been stolen, and in another, Person-2 makes inferences about what kind of large cat she is looking at while visiting a zoo. Further, we might say these quotidian beliefs correspond to cases of *ordinary incredulity*, not skepticism.[17] That is to say, Turri's vignettes do not evaluate a protagonist's beliefs in the realization of alternative skeptical hypotheses: that one is a brain in a vat, one is dreaming, and so forth. To the contrary, Turri's vignettes presuppose that such skeptical scenarios are *not realized.* There is indeed an external world such that Person-1 might worry whether his car is stolen, or Person-2 might visit the zoo and inquire into the nature of a feline. With further observation or sound inferential methods, the protagonists might alleviate their doubts. It is in this way Turri constructs a *strawman* against the skepticism he is after.

Skepticism is most powerfully argued for by using world-changing negative inferences: We are not in the world we think we are, but in a different world. We are some kind of thinking things, fooled by an evil demon, trapped in the Matrix (or some other computer simulation), always dreaming, or are unbodied brains floating in vats being given inputs from some machine. These world-changing inferences powerfully argue for skepticism because they all invoke an entire world strongly different from the one we think we inhabit, but a world that is *experientially invariant* from the world we think we inhabit. Turri discusses such world-changing inferences in the beginning of his paper, but he doesn't experiment on them at all. Here is an example of such a world-changing argument—we dub this argument *Hands*:

H1. If I know I have hands, then I know I am not a brain in a vat.
H2. But I don't know that I'm not a brain in a vat.
H3. Therefore, I don't know I have hands.

---

[17] See Klein (2015) for pertinent discussion.

H2 no doubt strikes the reader as true, or at least credibly true. But premise H2 gets its *credibility* not from being a part of a negative inferential belief (which it is), but from participating in a deep truth: the brain-in-a-vat scenario is experientially indistinguishable from our experiences in what we take to be the ordinary non-Matrix world. Perhaps Hands gets its force, appeal, and emotional power from our psychology, but it gets its *logical plausibility* as a correct description of the way things are from experiential invariance: if we were brains in vats, it would be impossible for us to discover this fact.[18]

To avoid an objection of equivocation, Turri would perhaps argue here that the psychological appeal of Hands and its logical plausibility are closely tied together. Such a move is plausible because there are cases in which an argument lacks psychological appeal to someone precisely because the argument lacks logical plausibility: One can become averse to a position one deems lacking in logical strength. However, an argument which enjoys psychological appeal from a human bias cannot *ipso facto* lack logical plausibility. Yet this is the very conclusion Turri draws. Recall Turri concludes, from his findings, that it is not "some deep fact about the nature of knowledge", but our psychology that invites skepticism (2015, p. 320). But the psychological appeal of a position does not *entail* the position's lack of logical strength. Turri is seeking to explain skepticism's emotional aspect and instead we are focusing on its logical aspects. The two aspects are distinct. And again, for Turri to insist that this distinction isn't real, or isn't real in this case, is to beg the question against the arguments presented here.

By focusing on quotidian beliefs, Turri has found that certain negative inferential beliefs have a kind of epistemically destructive appeal…. An interesting result. But this result hardly undermines the logical plausibility of an argument like Hands.

We conclude that Turri's experiments, while interesting and even useful, do not so much as dent skepticism.

4.5 Objection 3: Turri's Logical Bind

Turri has to hold that not all negative inferential beliefs are problematic, that not all of them are to be rejected because his own experiment-based argument involves negative inferential beliefs. We call this next argument *Turri's Argument* (*TA* for short).

TA1. If skepticism is philosophically tenable, then the negative inferential beliefs upon which skepticism depends are bias-free.
TA2. But the negative inferential beliefs upon which skepticism depends are not bias-free; indeed they are bias-laden.
TA3. Therefore, skepticism is not philosophically tenable.

---

[18] Putnam (1981) holds that if we were brains in vats we could not even think or wonder if we were brains in vats (due to his causal theory of meaning and reference). According to Putnam, if we think we might be brains in vats, we aren't. Not surprisingly, Putnam's argument is here rejected because the causal theory it requires is rejected.

Let's dispel immediately any worries of equivocation between the psychological appeal and philosophical tenability of the skeptical argument. This objection was discussed in the previous section. Instead let's focus here on the following move. Unless Turri is willing to banish *modus tollens* and then abandon his project, he must grant the negative inferential belief contained in his argument and its conclusion. We assume that Turri embraces *modus tollens*.

So, not all negative inferential beliefs are bad, on Turri's own view. Well, which ones are good? Clearly his in TA is. But we have argued in Sect. 4.4 (Objection 2), that other negative inferential beliefs, the world-changing ones, which depend on experiential invariance, are also good, and for logical reasons. But these latter beliefs result in skepticism. The only negative inferential beliefs to be rejected on the basis of Turri's experiments are the quotidian ones—and probably not even all of them, since negative inferential beliefs are ubiquitous and ineluctable.

Turri is now in the following logical bind:

> To maintain his experiment-based anti-skepticism project, he has to reject the argument Hands (and all similar arguments). To do this, he has to reject *all* negative inferential beliefs—the quotidian ones (which he rejects in his experiments), and the world-changing ones (since they fund skepticism), as well as other such beliefs like those supplied by *modus tollens*. But if he does this, he must abandon his experiment-based anti-skepticism project, which depends on drawing negative inferences, leading to negative inferential beliefs about skepticism.

No doubt Turri would claim to be able to avoid this bind by insisting only the negative inferential beliefs responsible for skepticism (along with the ones in his experiments) are to be rejected. We don't see any available argument for any such special pleading.

4.6 Skepticism

Three objections against Turri's anti-skepticism project have been examined. It seems clear now that Turri's claims that skepticism can be explained away as a mere function of our psychology aren't correct. We conclude that readers of Turri's paper (including Turri) cannot know that he has successfully refuted skepticism. At best, such readers should suspend their judgment—which is a kind of skepticism about Turri's project. By surviving Turri's new, experimental attack, skepticism, then, should look stronger than ever. Rather than explaining skepticism away, Turri has paved way for a new and more virulent skepticism.

5 Conclusion

This paper has examined three powerful clashes between science and philosophy. In particular, we have investigated three putative solutions: First to Zeno's paradoxes, then to the Frame Problem, and finally to skepticism. Examination of the three cases

has, in each case, revealed a limitation of scientific inquiry: to solve these three problems raised by philosophers, scientists must question-beggingly deploy premises and methods rejected by those very philosophers who constructed the problems. These deep and longstanding philosophical problems remain irresoluble and, so, implacable.[19]

References

Benardete J (1964) Infinity. Clarendon (Oxford University) Press, Oxford

Connell Louise, Lynott Dermot (2014) Principles of representation: why you can't represent the same concept twice. Top Cognit Sci 6:390–406

Dennett D (1978) Brainstorms: philosophical essays on mind and psychology. Bradford Books, Cambridge

DeRose K (2009) The case for contextualism: knowledge, skepticism, and context. Oxford University Press, New York

Dietrich E (2000) Analogy and conceptual change, or You can't step into the same mind twice. In: Dietrich E, Markman A (eds) Cognitive dynamics: conceptual change in humans and machines. Lawrence Erlbaum, Mahwah, pp 265–294

Dietrich E (2010) Analogical insight: toward unifying categorization and analogy. Cognit Proc 11(4):331

Dietrich E (2011) There is no progress in philosophy. In: Dietrich E, Weber Z (eds) Essays in philosophy, vol 12, issue date: July 2011, issue topic: philosophy's future: science or something else?

Dretske F (1970) Epistemic operators. J Philos 67:1007–1023

Dretske F (1971) Conclusive reasons. Aust J Philos 49:1–22

Fields C (2013) How humans solve the frame problem. J Exp Theor Artif Intell. https://doi.org/10.1080/0952813X.2012.741624

Fodor JA (1987) Modules, frames, fridgeons, sleeping dogs, and the music of the spheres. In: Pylyshyn (ed) The Robot's dilemma: the frame problem in artificial intelligence. Ablex, Norwood, pp 139–149

Fumerton R (1994) Skepticism and naturalistic epistemology. Midwest Stud Philos 19:321–340

Hayes P (1987) What the frame problem is and isn't. In: Pylyshyn Z (ed) The Robot's dilemma: the frame problem in artificial intelligence. Ablex Publishing, Norwood

Kaplan R, Kaplan E (2003) The art of the infinite. Oxford University Press, Oxford

Kelly T (2005) Moorean facts and belief revision, or can the skeptic win? Philos Perspect 19:179–209

Klein P (2015) Skepticism. The Stanford Encyclopedia of Philosophy. https://plato.stanford.edu/entries/skepticism/

Lewis D (1996) Elusive knowledge. Aust J Philos 74:549–567

McCarthy J, Hayes P (1969) Some philosophical problems from the standpoint of artificial intelligence. In: Meltzer B, Michie D (eds) Machine intelligence, vol 4. Edinburgh, Edinburgh University Press, pp 463–502

---

[19] I thank Chris Fields for comments on previous versions of this paper.

Moore GE (1962) Proof of an external world. Philosophical papers. Collier Brooks, New York, pp 144–148

Nozick R (1981) Philosophical explanations. Harvard University Press, Cambridge

Palmer J (2012) "Parmenides," The Stanford encyclopedia of philosophy, Summer 2012 Edition, Edward N. Zalta (ed). http://plato.stanford.edu/archives/sum2012/entries/parmenides/

Papa-Grimaldi A (1996) Why mathematical solutions of Zeno's Paradoxes miss the point: zeno's one and many relation and Parmenides' prohibition. Rev Metaphys 50:299–314

Pryor J (2000) The skeptic and the dogmatist. Noûs 34:517–549

Putnam H (1975) The meaning of 'Meaning'. In: Philosophical papers, vol 2: mind, language and reality. Cambridge University Press

Putnam H (1981) Brains in a vat. In: Reason, truth, and history, Chap. 1. Cambridge University Press, Cambridge, pp 1–21

Quine WVO (1975) The nature of natural knowledge. In: Guttenplan (ed) Mind and language. Clarendon Press

Robinson JM (1968) An introduction to early Greek philosophy. Houghton Mifflin Company, Boston

Russell B (1903) Principles of mathematics. Cambridge University Press, Cambridge

Salmon W (1975) Space, time, and motion: a philosophical introduction. Dickenson Press, Encino

Shanahan M (2016). The frame problem. In: The Stanford encyclopedia of philosophy (Spring 2016 Edition), Edward N. Zalta (ed). https://plato.stanford.edu/archives/spr2016/entries/frame-problem/

Turri J (2015) Skeptical appeal: the source-content bias. Cognit Sci 39:307–324