



# Merely superficially contingent a priori knowledge and the McKinsey paradox

Joshua Rowan Thorpe<sup>1</sup>

Received: 30 June 2021 / Accepted: 1 February 2022 / Published online: 24 February 2022  
© The Author(s) 2022

## Abstract

The conclusion of the McKinsey paradox is that certain contingent claims about the external world are knowable a priori. Almost all of the literature on the paradox assumes that this conclusion is unacceptable, and focuses on finding a way of avoiding it. However, there is no consensus that any of these responses work. In this paper I take a different approach, arguing that the conclusion is acceptable. First, I develop our understanding of what Evans calls *merely superficially contingent a priori knowledge*, and explain why there is no reason to deny that merely superficially contingent a priori knowledge of the external world is possible. I then argue that, properly understood, the conclusion of the McKinsey paradox is that merely superficially contingent knowledge of certain claims about the external world is possible, and so the conclusion is acceptable. Finally, I respond to the two main arguments that the conclusion of the paradox is unacceptable.

**Keywords** Epistemology · Philosophy of language · A priori · McKinsey paradox · Scepticism

## 1 Introduction

The McKinsey paradox is driven by two ideas.<sup>1</sup> The first is externalism about thought content, the view that some facts about mental content do not supervene on the facts about the intrinsic properties of the subject. The second is the thesis of privileged

---

<sup>1</sup>Originally put forward in McKinsey (1991), and reiterated in McKinsey (1994, 2007). Boghossian (1997) gives an influential formulation of the paradox.

<sup>2</sup>It may not be completely happy to characterise the thesis of privileged access in terms of a priori knowledge. Nonetheless, it has become standard to do so in discussions of the McKinsey paradox, and I continue the practice here.

---

✉ Joshua Rowan Thorpe  
joshthop@gmail.com

<sup>1</sup> School of Philosophy, Psychology and Language Sciences, University of Edinburgh, 3 Charles Street, Edinburgh, Scotland EH89AD, UK

access, according to which we sometimes have a priori knowledge of what we are thinking.<sup>2</sup> Each of these ideas is independently plausible, at least upon suitable refinement. Taken together, however, they seem to have the consequence that certain contingent claims about the external world are knowable a priori. This conclusion is widely thought to be unacceptable. Hence the paradox.

Almost all responses to the paradox have attempted to avoid the allegedly unacceptable conclusion. Despite a great deal of discussion, however, there is no consensus that any of these responses work. In this paper I take a different approach, arguing that the conclusion of the paradox is *acceptable*. By this I do not mean that the conclusion is true, but that there is no good reason to deny it. So, if commitment to externalism and the thesis of privileged access does indeed put pressure on us to accept the conclusion of the paradox, we can simply accept it.<sup>3</sup>

I make this argument by developing our understanding of what Evans (1979) calls *merely superficially contingent a priori knowledge* (henceforth SCAP knowledge).<sup>4</sup> SCAP knowledge is a priori knowledge of a contingent thought that is, nevertheless, semantically guaranteed. Having improved on Evans' definition of a semantic guarantee, I explain why there is no reason to deny that SCAP knowledge of the external world is possible. I then show that, properly understood, the conclusion of the McKinsey paradox is that I can have SCAP knowledge of certain claims about the external world. This insight provides positive reason to think that the conclusion of the paradox is acceptable.

I then consider potential reasons to think that the conclusion of the paradox is unacceptable. Two such reasons have been offered in the literature: it has been suggested that the conclusion offers too quick a response to scepticism, and it has been suggested that the conclusion is counterintuitive. These arguments depend on the implicit assumption that the conclusion of the paradox is that it is possible for a subject to have *deeply contingent a priori knowledge*—a priori knowledge of a thought that is contingent and *not* semantically guaranteed—of certain claims about the external world. This assumption is false, so both arguments fail.

## 2 The paradox

The McKinsey paradox is typically presented as a paradox about what can be known a priori, where a priori knowledge is understood as knowledge whose justification does not include sensory evidence. According to a simple version of the paradox, I can engage in the following a priori reasoning:

<sup>3</sup> To my knowledge, the only other person to take this approach is Sawyer (1998). I highlight some of the differences between mine and Sawyer's responses in fn. 22 below. Pryor (2007) cites Putnam (1981) and his defenders as responding to the paradox by embracing its conclusion, but this is implausible. Putnam was responding to scepticism, and the McKinsey paradox was not in print at the time. Moreover, I think that there are important differences between Putnam's response to scepticism and the argument of the McKinsey paradox. However, I don't have space to argue this here.

<sup>4</sup> The 'M' is dropped from the acronym to make it pronounceable.

- (1) I am thinking that water is wet, and
- (2) If I am thinking that water is wet, I am in a water environment; so
- (3) I am in a water environment.<sup>5</sup>

Here, a ‘water environment’ is just an environment that does or did contain water.<sup>6</sup>

The claim that I can know (1) a priori is supported by the thesis of privileged access, according to which subjects can sometimes have a priori knowledge of what they are thinking. Of course, not all of our knowledge of what we think is a priori. We sometimes discover that we have a certain thought by observing our own behaviour, or by being told. Nonetheless, it does not seem that I typically need to appeal to my sensory evidence to justify my belief that I am currently thinking that water is wet.

The claim that I can know (2) a priori is supported by the idea that I can have a priori knowledge that a particular form of externalism about thought content is true. According to this view, which we can call *water externalism*, causal contact with water necessarily plays a role in fixing the content of the concept ‘water’. Thus, to have the concept ‘water’ a subject must be in a water environment. Water externalism is typically supported by a priori reflection on thought experiments, such as Putnam’s Twin Earth thought experiment.<sup>7</sup>

In fact, the considered externalist view of the concept ‘water’ is likely to be more complicated than water externalism as just described. Perhaps, for example, to have the concept of water one must have been in causal contact with water, or with experts who have theorised about the existence of water, or have theorised about water oneself.<sup>8</sup> Perhaps there will be further disjuncts. However, these complexities are unlikely to affect our discussion: a more sophisticated version of externalism is likely to complicate but not block the paradox,<sup>9</sup> and it would not affect the response to the paradox that I ultimately want to give. We can keep things simple by focusing on the version of the paradox generated by water externalism from here on.<sup>10</sup>

Almost all of the literature on the Mckinsey paradox assumes that the conclusion that I can know a priori that I am in a water environment is unacceptable, and attempts to solve the paradox have focused on finding ways to avoid this conclusion. For example, it has been suggested that we:

<sup>5</sup> The conclusion of this reasoning is, effectively, that the concept ‘water’ is not empty. This makes the conclusion particularly striking, but it is not integral to the paradox. See for example versions of the paradox discussed in Brown (1995) and Sawyer (1998).

<sup>6</sup> Sawyer (1998) also uses the term ‘water environment’, but in order to take account of the more sophisticated externalist view that I mention below her definition includes environments in which there are or were experts who have theorised about water.

<sup>7</sup> Putnam (1973, 1975) explicitly intends the Twin Earth scenario to support externalism about linguistic meaning. McGinn (1977) and Burge (1982) adapt it to support externalism about thought content.

<sup>8</sup> See Burge (1979) for this sort of view.

<sup>9</sup> As Brown (1995) argues.

<sup>10</sup> It has been argued that if water externalism is true then my knowledge that I am thinking that water is wet would rest on my empirical knowledge that I am in a water environment, and thus that it would not be a priori; see for example Boghossian (1989). See Falvey and Owens (1994) and Gertler (2015) for convincing replies.

- Give up on externalism altogether, thus denying (2).<sup>11</sup>
- Deny that we can know a priori that any particular concept is externalist, thus denying that (2) is knowable a priori.<sup>12</sup>
- Modify the thesis of privileged access, thus denying that (1) is knowable a priori.<sup>13</sup>
- Accept that I can have a priori knowledge of (1) and (2), but deny that I can reason from this knowledge to attain a priori knowledge of (3).<sup>14</sup>

Each solution has its supporters, but there is no consensus that any of them work.<sup>15</sup>

Somewhat surprisingly given the focus on avoiding the conclusion that I can know (3) a priori, there is very little explanation in the literature on the paradox of why the conclusion that I can know (3) a priori is unacceptable.<sup>16</sup> Given the energy that philosophers have expended on avoiding it, it is worth asking whether this conclusion really is unacceptable. I will argue that it is not. If endorsement of externalism and privileged access really do put pressure on us to accept this conclusion, there is no cost to accepting it. That is my solution to the paradox.

### 3 Merely superficially contingent a priori knowledge

In this section I explain the distinction between SCAP knowledge and deeply contingent a priori knowledge. I then argue that the reasons for thinking that deeply contingent a priori knowledge is impossible do not apply to SCAP knowledge. Indeed, there is no good reason to think that SCAP knowledge is impossible.

It is not a conceptual truth that if a proposition is a priori it is necessary.<sup>17</sup> And in fact, there are numerous potential examples of contingent a priori knowledge of the external world.<sup>18</sup> Consider the following potential example of contingent a priori knowledge adapted from Evans (1979).<sup>19</sup> Suppose that I stipulate that the reference of 'Julius' is to be fixed by the description: the unique inventor of the zip at the actual world.<sup>20</sup> Having done so, I can see that my thought that.

<sup>11</sup> McKinsey (1991, 2007).

<sup>12</sup> Gallois and Hawthorne (1996).

<sup>13</sup> Pryor (2007).

<sup>14</sup> Davies (1998) and Wright (2000, 2011) argue that the argument from (1) and (2) to (3) is unsuccessful due to a warrant transmission failure. Sawyer (2015) argues that once we appreciate that self-knowledge is 'contrastive' we will see that this argument is invalid.

<sup>15</sup> Kallestrup (2011) gives an overview of solutions to the paradox, and their problems.

<sup>16</sup> The two arguments that this conclusion is unacceptable discussed in §5 are both extrapolations from a few suggestive comments.

<sup>17</sup> As Kripke (1980) points out.

<sup>18</sup> To avoid getting bogged down in details I'll keep my presentation of these examples fairly intuitive, since nothing turns on whether we accept any particular one of them.

<sup>19</sup> Evans' original example concerns sentences, rather than thoughts. However, since Evans clearly thought that it was possible to have knowledge of what the sentences in his examples meant, this difference doesn't cut all that deep.

<sup>20</sup> See Ruffino (2020) for a compelling argument that certain central examples of the contingent a priori are to be understood as the result of stipulative speech acts.

if anyone uniquely invented the zip it was Julius.

is true if.

if anyone uniquely invented the zip, the unique inventor of the zip invented the zip.

The latter is a logical truth, knowable a priori. Thus, I can infer, and therefore come to know a priori that if anyone uniquely invented the zip, it was Julius. And yet, it is contingent that if anyone uniquely invented the zip it was Julius. There are possible worlds where the antecedent of this conditional is true, and yet the consequent is false, namely, worlds in which someone uniquely invented the zip, but it wasn't the person rigidly designated by the name 'Julius' in this world.<sup>21</sup> So my knowledge that if anyone uniquely invented the zip it was Julius seems to be contingent a priori knowledge. Let us call this the Julius example.

Evans classifies the Julius example as a case of merely superficially contingent a priori knowledge (SCAP knowledge), a priori knowledge of a contingent claim that is, nevertheless, semantically guaranteed. For Evans, a thought is semantically guaranteed when a priori knowledge of the meaning of the thought, in conjunction with background a priori knowledge, suffices for knowledge that the thought is true. However, this characterisation of semantic guarantee does not seem to fit the Julius example. In this example, what does the work is not my a priori knowledge of the meaning of my thought, but rather my a priori knowledge of *how the content of the thought was determined*. In particular, it is my a priori knowledge that I stipulated that 'Julius' refers to the unique inventor of the zip at the actual world that allows me to know a priori that if anyone uniquely invented the zip it was Julius. To take account of this we will alter Evans' definition of a semantic guarantee and say that a thought is semantically guaranteed when a priori knowledge of how the content of the thought was determined, in conjunction with background a priori knowledge, suffices for knowledge that the thought in question is true.

This definition of semantic guarantee fits with other potential examples of SCAP knowledge. To adapt an example from Kripke (1980), let us say that I stipulate the meaning of 'one metre\*' by pointing at a stick, S, and saying: that is one metre\* long.<sup>22</sup> One might think that I can know that S is one metre\* long a priori, merely by reflecting on the way in which I fixed the meaning of 'one metre\*'. Call this the Stick example.

One might worry that my knowledge that S is one metre\* will include empirical knowledge that S exists in its justification, and thus will turn out to be empirical knowledge after all. Whether or not this worry is correct, we can manufacture a similar example that is not subject to it. I stipulate the meaning of 'F' by saying: a state of affairs is F if and only if it is exactly the same as the total state of affairs at this world. I then have a priori knowledge of the contingent claim that the total state of affairs is F. If there is an external world at all, then this will be a priori knowledge of a contingent claim about the external world. Call this the World example.

<sup>21</sup> It is worth noting that this example depends on Evans' claim that 'Julius' is a rigid designator—a claim that Evans offers no argument for beyond our intuitions as speakers.

<sup>22</sup> Kripke's original example involved the standard metre in Paris.

Why think that the claim that we can have SCAP knowledge acceptable? To answer this question, we need to know why the claim that we can have *deeply* contingent a priori knowledge—a priori knowledge of thoughts that are not semantically guaranteed—is unacceptable. So far as I can see, there are two *prima facie* reasons to think that deeply contingent a priori knowledge is unacceptable.<sup>23</sup>

First, when Evans (1979, p. 179) claims that the idea that we can have deeply contingent a priori knowledge is ‘intolerable’, his thought seems to be that there can be no explanation of how we might come to know that a contingent thought with no semantic guarantee is true, except by appealing to our sensory experiences. Thus, it is *inexplicable* how we might have deeply contingent a priori knowledge.

The second reason, not noted by Evans, stems from the suspicion that if a deal looks too good to be true, it probably is. To get contingent knowledge of the external world, one pays the price of engaging in empirical enquiry. The price does vary: sometimes the amount of empirical enquiry required is more, sometimes less. But deeply contingent a priori knowledge is meant to be knowledge of the external world without paying the cost of *any* epistemic enquiry. Doesn’t that sound like a deal that’s too good to be true?

However, neither of these reasons applies to SCAP knowledge. Regarding the first reason, deeply contingent a priori knowledge may be inexplicable, but SCAP knowledge is not. One can know a merely superficially contingent proposition by noticing that the proposition is semantically guaranteed to be true. Thus, there is no mystery as to how one can know such a proposition without appealing to sensory evidence. This appears to be the reason that Evans found the idea of SCAP knowledge acceptable.<sup>24</sup>

Regarding the second reason, sometimes when a deal looks too good to be true it turns out that the explanation is that you can indeed make the deal, but that what you get isn’t as valuable as you thought it was going to be. This is the case with SCAP knowledge, due to an important feature of SCAP knowledge that Evans appears to have missed. Namely, that SCAP knowledge is not as valuable to the knower as it may at first appear to be, due to the fact that it is, as I will say, *a prioristically inferentially inert* (henceforth APII). Knowledge is APII when, so long as I am confined to what I can know a priori, I cannot make most of the interesting deductions that I would usually be able to make from it.

To illustrate, consider the Stick example in which I know that S is one metre\* long. So long as I am confined to what I can know a priori, I cannot make many interesting or useful deductions from this knowledge. For example, I cannot deduce that S is long, that S is short, or that S is/is not one hundred and fifty (normal) centimetres.

<sup>23</sup> It has been suggested that the claim that we can have deeply contingent a priori knowledge is not unacceptable; see for example Williamson (1989) and Hawthorne (2002). I won’t argue against this suggestion here, but note that those who are bothered by the conclusion of the McKinsey paradox are unlikely to be sympathetic to it. It would be odd to think that examples of deeply contingent a priori knowledge are many and various, but to be troubled by the example of the conclusion of the McKinsey paradox. At least, someone who took this position would need to say something about why the conclusion of the paradox is so much more troubling than other examples of deeply contingent a priori knowledge, but so far as I know no one has attempted to do this.

<sup>24</sup> Evans is not explicit about why he thinks that merely superficially contingent knowledge is acceptable while deeply contingent a priori knowledge is not, but this seems to be what he has in mind. Hawthorne (2002, pp. 247–248) reads Evans in this way.

This is because I cannot have a priori knowledge of any specification of the length to which my newly defined term ‘one metre\*’ refers that would allow me to make these inferences. I don’t know, for example, that ‘one metre\*’ refers to a fairly short length, or to a length that is one hundred and fifty centimetres long.

It is true that I can make *some* a priori deductions from my knowledge that S is one metre\*, because I do have some a priori knowledge of what ‘one metre\*’ refers to. Since it uses a term to state its own reference, there is no need for empirical investigation to verify the disquotational claim that ‘one metre\*’ refers to lengths of one metre\*. This knowledge of the referent of ‘one metre\*’ allows me to make relatively trivial deductions from my knowledge that S is one metre\* even while confined to what I can know a priori. For example, I can deduce that S is one metre\*, that S is twice the length of half of one metre\*, and that S is one metre\* or the moon is made of cheese. But these trivial inferences are not interesting to me in the way that the inference that S is short would be.

There will however be a good sense in which what I know has interesting *consequences*. Suppose that in fact one metre\* refers to lengths of one hundred and fifty centimetres. In that case, the fact that S is one metre\* long will, at the very least, metaphysically necessitate that S is one hundred and fifty centimetres long, and that S is fairly short. One may wish to say that my knowledge that S is one metre\* long entails that S is one hundred and fifty centimetres long, depending on how exactly one thinks of propositions and their entailment relations. For our purposes, it does not matter how exactly we think of the interesting consequences of what I know. The point of saying that my knowledge that S is one metre\* is APII is that, so long as I am confined to what I can know a priori, I am blind to most of these consequences.

The recognition that my knowledge that S is one metre\* long is APII helps to dispel any feeling that I am illicitly acquiring an epistemic good that I should only be able to get at the cost of empirical enquiry. Perhaps the idea that I can get useful knowledge of the length of S that would allow me to make interesting inferences about S without engaging in empirical work is too good to be true. But it is not so surprising if I can get a priori knowledge of the length of S that, so long as I am confined to what I can know a priori, is mostly inferentially useless. When it comes to SCAP knowledge, you get what you pay for.

The SCAP knowledge in the two other examples is also APII. In the World example, so long as I am confined to what I can know a priori my only knowledge of the referent of ‘F’ is that it refers to states of affairs that are F, and that these states of affairs are exactly the same as the state of affairs at this world. So, although my knowledge that the state of affairs is F has many extremely specific and interesting consequences, I am blind to these consequences, so long as I am confined to what I can know a priori.

The Julius case is a little different. I can have a priori knowledge of the usual triviality that ‘Julius’ refers to Julius. However, since I also know a priori that ‘Julius’ refers to the unique inventor of the zip, I can deduce a priori that *if* Julius exists then they are an inventor, and (if it is a priori that inventors are sentient) that *if* Julius exists then they are sentient. This is a bit more substantial than the a priori deductions that can be made from previous examples of APII knowledge. However, what I can deduce is still severely limited. Since I can’t know a priori that there is a unique inventor of

the zip, there is no a priori way for me to discharge the antecedents of these bits of conditional knowledge and thereby come to know that a sentient inventor exists.

There is good reason to think that SCAP knowledge will always be APII. The semantic guarantee in the Stick example is generated by the fact that I fixed the content of ‘one metre\*’ by referring to the stick, S. However, suppose that I attempt to stipulate further content of the concept ‘one metre\*’ by saying that it refers to one hundred and fifty centimetres. If I could successfully add this extra content while maintaining the semantic guarantee, I could make many interesting a priori deductions from my knowledge that S is one metre\*. However, I cannot do this, because there is no guarantee that the new stipulation is compatible with my original ostensive stipulation. I may get around this by saying that if the stipulations are incompatible the stipulation that the content of ‘one metre\*’ is to be fixed by the length of S is to take precedence in fixing the reference. In that case the semantic guarantee of my thought that S is one metre\* long would be preserved, but my ability to make the extra interesting a priori deductions from it would be lost. Or, I could say that the stipulation that ‘one metre\*’ is to refer to lengths of one hundred and fifty centimetres takes precedence in fixing the reference. In that case I could make interesting a priori deductions from my thought that S is one metre\*, but this thought would no longer be semantically guaranteed. The lesson here is that if it is to provide a semantic guarantee, the content fixing stipulation which potentially generates this guarantee must take precedence in fixing the reference over any other potentially incompatible stipulations of content that might allow interesting a priori inferences. This ensures that genuinely semantically guaranteed thoughts are APII.

#### 4 Solving the paradox

To refresh: the McKinsey paradox arises because commitment the thesis of privileged access and the a priori knowability of water externalism seem to allow me to engage in the following a priori reasoning:

- (1) I am thinking that water is wet, and
- (2) If I am thinking that water is wet, I am in a water environment; so
- (3) I am in a water environment

We are now in a position to see that if I can indeed engage in this reasoning, the resulting a priori knowledge of (3) is SCAP knowledge.

The paradox requires that my a priori knowledge of (2) is underwritten by my a priori knowledge that water externalism is true. If I have a priori knowledge of water externalism and its implications, then my thought that I am in a water environment is semantically guaranteed, because water externalism entails that the content of the concept ‘water’ was fixed by my coming into causal contact with water. As a result, I need only reflect on how the content of my thought that I am in a water environment was fixed to see that this thought is true. Thus, my a priori knowledge that I am in a water environment is, like my knowledge in the Julius, Stick and World examples, SCAP knowledge.

My a priori knowledge that I am in a water environment and the examples of SCAP knowledge given in §3 are not the same in all respects. One important difference is that in the examples of §3 the semantic guarantee of the relevant thoughts arises due to my a priori knowledge that I have performed an act of stipulation to determine the content of the relevant thought, thereby guaranteeing that the thought is true. This act of stipulation is a thought—for example, in the Julius case I think: let the reference of ‘Julius’ be the unique inventor of the zip at the actual world. I know a priori that I had this thought, and therefore what the act of stipulation was and that it took place. By contrast, water externalism says that the content of the concept ‘water’ is fixed by causal contact with water, rather than by an act of stipulation. My knowledge of the truth of water externalism is a priori (or so we must suppose, if we are to generate the McKinsey paradox), and so my knowledge that the content of my thought that I am in a water environment is determined in this way does not come from a priori introspection, but from a priori knowledge of philosophical theory.

However, this difference does not undermine the claim that my a priori knowledge that I am in a water environment is like my a priori knowledge that if anyone uniquely invented the zip it was Julius *in that both instances of knowledge are underpinned by a semantic guarantee*. All that this claim requires is that both instances of a priori knowledge are underpinned by a semantic guarantee: a priori knowledge that the content of the thought was determined in such a way as to guarantee that the thought is true. And in both cases, the relevant thoughts are semantically guaranteed, even if the source of the semantic guarantee is different: a priori introspection in the Julius case, and a priori knowledge of philosophical theory in the case where I work through the reasoning of the McKinsey paradox.

Since it is SCAP knowledge, we would expect my knowledge that I am in a water environment to be APII. And it is: I can know a priori that the referent of ‘water’ is water, and that it is something with which I have had a certain sort of causal interaction, but not much more. Thus, I can make few interesting a priori deductions from what I know. I cannot deduce a priori that I am in an environment that does or did have something drinkable or wet in it, for example.

Since the conclusion of the paradox is that I can have SCAP knowledge of the external world it is acceptable, for the same reason that the claim that I can have SCAP knowledge in the examples considered in §3 is acceptable. First, we can explain how I can know a priori that I am in a water environment by referring to that thought’s semantic guarantee.<sup>25</sup> Second, my a priori knowledge that I am in a water environment is not as interesting as it might have seemed, due to being APII (I’ll say more about this second point in §5).

I want to end this section by considering a worry about the line of thought that we have been pursuing. It is doubtful, at least on any plausible externalist view, that

<sup>25</sup> This part of my response is very much in keeping with Sawyer (1998), who responds to the McKinsey paradox by arguing that the idea that I can have a priori knowledge that I am in a water environment due to externalism and privileged access is acceptable because it is explicable. However, Sawyer does not explain this knowledge in terms of a semantic guarantee. This deprives her of the comparison with other potential examples of SCAP knowledge, and of the point that my knowledge that I am in a water environment is APII. (Sawyer (2015) went on to change her mind, taking up the common assumption that the conclusion of the paradox is unacceptable).

there is any *particular* bit of non-trivial knowledge about the referent of a concept that one must have to possess that concept. However, one might still insist that one must have *some* non-trivial knowledge of the referent of a concept that one possesses. If this is so, then one might think that A<sub>PII</sub> thoughts are impossible, because they are characterised precisely by lack of non-trivial knowledge of the referents of some of the concepts such thoughts involve.

To think this would be a mistake, however. A<sub>PII</sub> thoughts are characterised by lack of non-trivial a priori knowledge of the referents of some of the concepts that they involve. Lack of non-trivial a priori knowledge of the referent of e.g. ‘water’ is compatible with possession of *empirical* non-trivial knowledge of the referent of ‘water’. If it is a necessary condition of my possession of the concept ‘water’ that I have some substantial empirical knowledge of its referent, then this empirical knowledge is a necessary condition of my thinking that I am in a water environment. However, this is not incompatible with the possibility of my having SCAP knowledge that I am in a water environment. It is commonly recognised that although a priori knowledge cannot include empirical knowledge in its justification, it may nevertheless depend on empirical knowledge in other ways.<sup>26</sup> Of course, if I bring my empirical knowledge of the referent of ‘water’ together with my SCAP knowledge that I am in a water environment, I will be able to make interesting deductions from this knowledge, but it isn’t surprising that combining a priori and empirical knowledge allows me to make interesting deductions about my environment.

## 5 Why not?

Let us turn to the two main arguments against the acceptability of the conclusion of the McKinsey paradox. It has been suggested that the conclusion is too quick as a response to scepticism, and that it is counterintuitive.<sup>27</sup> Both of these arguments fail, because they implicitly rely on the false assumption that the conclusion of the paradox is that we can have deeply contingent a priori knowledge.

Sawyer (2015, p. 75) suggests that the conclusion of the paradox is unacceptable because it ‘seems to provide an implausibly quick response to the external world sceptic.’<sup>28</sup> Unfortunately, she does not expand on what she means by external world scepticism, or explain why the conclusion of the paradox would refute it. In fact, the conclusion of the paradox does not refute any interesting version of external world scepticism.

One familiar form of scepticism starts by asking me to consider a scenario such as the following:

Vat Scenario: I always have been and always will be a brain-in-a-vat. A supercomputer feeds me sensory experiences that suggest that everything is as I ordinarily

<sup>26</sup> See for example Jenkins and Kasaki (2015).

<sup>27</sup> It is worth noting that it is difficult to find in the literature on the McKinsey paradox anything approaching a full paragraph on why the conclusion of the paradox is unacceptable. As a result, my description of these arguments necessarily involves a good deal of extrapolation.

<sup>28</sup> Gallois and Hawthorne (1996) say something similar.

take it to be. However, aside from my brain in its vat and the supercomputer, nothing else exists.

The sceptic then argues that I do not know that I am not in the Vat Scenario. If the sceptic can establish that I do not know that I am not in the Vat Scenario, it seems to follow that I have very little (if any) empirical knowledge.<sup>29</sup> Let us call scepticism motivated by the claim that I do not know that I am not in the Vat Scenario *vat scepticism*.

Now it might seem obvious that if, as suggested by the McKinsey paradox, I can have a priori knowledge that I am in a water environment, I will be able to rule out the Vat Scenario on purely a priori grounds.<sup>30</sup> The claim that I am in an environment which does or did contain water is incompatible with my being in the Vat Scenario, so can't I just deduce from my a priori knowledge that I am in a water environment that I am not in the Vat Scenario, thus refuting vat scepticism?<sup>31</sup>

No. My SCAP knowledge that I am in a water environment is APII, and one manifestation of this is that I cannot deduce from it that I am not in the Vat Scenario. To elaborate: so long as I am confined to what I can know a priori, I cannot deduce from my knowledge that I am in a water environment that I am not in the Vat Scenario, because I cannot know a priori that my concept 'water' does not refer to something that would be in my environment in the vat scenario, such as, for example, certain electrical impulses that stimulate my brain, or the features of the computer program responsible for sending those electrical impulses.<sup>32</sup> To establish that my concept 'water' does not refer to any of these things, I would have to engage in empirical enquiry, just as I must engage in empirical enquiry to establish that my concept 'water' does not refer to, say, XYZ. Since I am concerned with scepticism regarding my empirical knowledge, it would be question begging to assume that I have such knowledge. However, without this knowledge of the referent of 'water', I cannot recognise that my knowledge that I am in a water environment has the consequence that I am not in the Vat Scenario.<sup>33</sup>

Of course, my knowledge that I am in a water environment does in fact have this consequence. At the very least, my being in a water environment metaphysically necessitates that I am not in the vat scenario, and, depending on how one thinks of propositions, one might also say that what I know entails that I am not in the vat scenario. The point here is that insofar as I am confined to what I can know a priori, I am blind to these consequences.

We may conclude that the conclusion of the McKinsey paradox cannot be used to refute vat scepticism, which is perhaps the most discussed form of scepticism in

<sup>29</sup> Cohen (1998) compares two reasons for thinking that this follows.

<sup>30</sup> This is sometimes taken to be the main thrust of Putnam's (1981) argument that I am not a brain in a vat.

<sup>31</sup> One might object that I *would* be in a water environment if I were in the Vat Scenario, since there would be some water in my brain, and possibly in the vat too. However, we could easily avoid this objection by switching to an externalist concept of something that would not exist at all if the Vat Scenario were true, e.g. the concept of a beech tree.

<sup>32</sup> See Putnam (1981, p. 14) on what a brain in a vat would have concepts of.

<sup>33</sup> The point made in this paragraph is similar to one made by Falvey and Owens (1994) in their criticism of Putnam's vat argument. In Thorpe (2019) and Thorpe and Wright (forthcoming) argue that their criticism of the vat argument does not touch the most plausible version of the vat argument, though I think that it applies to the anti-sceptical argument considered here.

contemporary epistemology. The impression that the conclusion of the paradox does refute vat scepticism stems from the mistaken idea that it implies that I have deeply contingent a priori knowledge that I am in a water environment, in which case there would be no reason to suppose that my knowledge that I am in a water environment is APII. If this were so, I might be able to recognise a priori that my thought that I am in a water environment entails that I am not in the Vat Scenario. But it isn't so.

One might still think that the conclusion of the paradox at least refutes a common formulation of scepticism according to which I can have no contingent knowledge of the external world, since it says that I have SCAP knowledge that I am in a water environment. Indeed, if we adopt a view of other concepts analogous to the water externalist view of the concept 'water', I will be able to give a similar a priori proof that my environment does or did contain the things to which those concepts refer as well.

However, we are now in a position to see that this is not an interesting formulation of the sceptical claim, since it would be refuted by any example of SCAP knowledge of the external world. But, surely, it was never important to the sceptic to deny that I can know that if anyone invented the zip it was Julius, or that the state of affairs is F. As we have seen, these bits of knowledge are not as interesting as they may at first seem, due to being APII. The sceptic wants to deny that I have any *interesting* knowledge of the external world, and she can easily reformulate her claim to reflect this by saying that I do not have any contingent knowledge of the external world that is not APII. The conclusion of the paradox does not refute this sceptical claim.

It might be protested that the conclusion of the McKinsey paradox at least offers a refutation of the classic sceptical claim that I do not know that there are any physical objects. I may not know whether 'water' refers to H<sub>2</sub>O or XYZ or features of a computer program, but at least I know that there is or was such a thing, and that I was in some sort of causal contact with it.

I'm not sure that it would really be cause for concern if the conclusion of the McKinsey paradox allows me to refute this form of scepticism. As Brueckner (1992, p. 118) says, although the idea that one can have a priori knowledge that includes 'much detail concerning the character of the external world' may be unacceptable, it is not obvious that a priori knowledge that 'is simply knowledge that something or other physical exists distinct from oneself' is impossible.<sup>34</sup> Brueckner points out that to deny the possibility of the latter knowledge would be to dismiss in advance the possibility of a successful Kant-style transcendental argument against scepticism. Although careful reflection may show that no such transcendental argument is successful, the unargued assumption that it is not is surely not a good basis for rejection of the conclusion of the paradox. Allowing the possibility of APII knowledge that I am in a water environment seems to be entirely in the spirit of Brueckner's assertion that the idea that a priori knowledge that something external to oneself exists is acceptable so long as it does not say anything about the detailed characteristics of the external world.

Putting Brueckner's point aside, however, I am not convinced that my knowledge that I am in a water environment would allow me to refute the sceptical claim that

---

<sup>34</sup> Philosophers who reject the conclusion of the paradox have allowed versions of Brueckner's point; see for example Brown (1995), and Ball (2007).

no physical objects exist, because I am not convinced that I can know a priori that my concept ‘water’ refers to physical objects, rather than, say actual and possible arrangements of my sensory experiences, or to Berkeleyan ideas. All I know a priori is that my concept ‘water’ refers to whatever bears the appropriate reference fixing causal relation to it. Just as I cannot know a priori that the thing at the other end of this causal chain is not XYZ, or a feature of a computer program, or a collection of electrical signals, I cannot know that it is not a collection of sensory experiences, or Berkeleyan ideas. Thus, I cannot deduce a priori from my SCAP knowledge that I am in a water environment that at least one physical object exists.

Let us move on to the idea that the conclusion of the paradox is counterintuitive. Having stated the McKinsey paradox, Pryor (2007, p. 179) claims that ‘it’s counterintuitive that you should be able to tell what your environment is like just on the basis of this kind of [a priori] reasoning’.<sup>35</sup> Putting worries about whether we have reliable intuitions about terms of art like ‘a priori knowledge’ aside, there are two problems with this claim.

The first is that it is not clear that you count as knowing what your environment is like because you have APII knowledge of your environment. Most people would not count me as knowing where I am if I were to announce that I know that I am here, despite having no interesting knowledge of the referent of ‘here’, and thus no ability to make any useful deductions from what I know. For the same reason, it is natural not to count me as knowing what my environment is like if, for example, I have APII knowledge that the total state of affairs is F. So, even if Pryor is right that it’s counterintuitive to say that you can tell what your environment is like on the basis of a priori reasoning, it is not clear that I count as knowing what my environment is like when this reasoning results in SCAP knowledge that I am in a water environment.<sup>36</sup>

The second problem arises from the fact that intuitions tend to subside or to lose their evidential force when we admit that we don’t have the same intuitions about cases that seem to be the same in all relevant respects. Thus, someone who wants to maintain that their intuition about the unacceptability of the conclusion that I can have SCAP knowledge that I am in a water environment is evidence against it faces a choice. Either, they can say that all other potential examples of SCAP knowledge are also intuitively unacceptable, or they can attempt to point out a relevant difference between potential examples of the merely superficially contingent a priori that are not intuitively unacceptable and the conclusion of the McKinsey paradox.

Some philosophers will not, on reflection, think that the potential examples of SCAP knowledge discussed in §3 are genuine, and some may eventually conclude that there are no genuine examples of SCAP knowledge at all, but I do not know of anyone who has claimed that these examples are intuitively unacceptable. On the other hand, it isn’t clear what the significant difference between the conclusion of the McKinsey paradox and these examples might be. I expect the most popular candidate to be that that the conclusion of the paradox allows me to deduce the existence of something in the external world, while other instances of SCAP knowledge do not. As we have seen, there is good reason to think that this is not the case. However, even putting this

<sup>35</sup> Davies (1998) also says that the conclusion of the paradox is counterintuitive.

<sup>36</sup> I am grateful to Bernhard Salow for this point.

point aside, it is not clear why this putative difference would preserve the force of the intuition that the conclusion of the paradox is unacceptable. As has already been noted, the idea of an a priori proof that *something* external to myself exists doesn't seem to be obviously beyond the pale, and various philosophers have taken themselves to have given such a proof.

## 6 Conclusion

The solution to the McKinsey paradox that I have proposed is that the conclusion of the paradox is acceptable. If endorsement of privileged access and externalism about thought content does put pressure on us to accept the possibility of a priori knowledge of certain contingent propositions about the external world, we can simply accept this consequence. This is because this knowledge would be SCAP knowledge, and there is no good reason to think that SCAP knowledge is impossible. Having recognised this, we can see that the two main arguments that this consequence is unacceptable fail.

For those who do not accept my conclusion, I hope that this paper will stimulate them to clarify *why* they think that the conclusion of the paradox is unacceptable. Am I wrong in claiming that the conclusion of the paradox is that we have SCAP knowledge? Or is there a reason to think that SCAP knowledge is impossible after all? In a debate that has been almost entirely motivated by the aim of avoiding the conclusion of the paradox, such clarification can only be a good thing.

**Acknowledgements** Thanks to Bernhard Salow, Crispin Wright, David Horst and Marco Ruffino for helpful discussion. I am also grateful to two anonymous *Synthese* referees for helpful comments.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Ball, D. (2007). Twin-earth externalism and concept possession. *Australasian Journal of Philosophy*, 83(3), 457–472.
- Boghossian, P. A. (1989). Content and self knowledge. *Philosophical Topics*, 17, 5–26.
- Boghossian, P. A. (1997). What the externalist can know a priori. *Proceedings of the Aristotelian Society*, 97, 161–75.
- Brown, J. (1995). The incompatibility of anti-individualism and privileged access. *Analysis*, 55(3), 149–156.
- Brueckner, A. (1992). What an anti-individualist knows a priori. *Analysis*, 52, 111–118.
- Burge, T. (1979). Individualism and the mental. In D. M. Rosenthal (Ed.), *The nature of mind* (Vol. 4, pp. 536–567). New York: Oxford University Press.
- Burge, T. (1982). Other bodies. In A. Woodfield (Ed.), *Thought and object* (pp. 97–120). Oxford University Press.

- Cohen, S. (1998). Two kinds of skeptical argument. *Philosophy and Phenomenological Research*, 58(1), 143–159.
- Davies, M. (1998). Externalism, architecturalism, and epistemic warrant. In C. Wright, M. Smith, & C. Macdonald (Eds.), *Knowing our own minds: Essays in self-knowledge* (pp. 321–361). Oxford University Press.
- Evans, G. (1979). Reference and contingency. *The Monist*, 62(2), 161–189.
- Falvey, K., & Owens, J. (1994). Externalism, self knowledge, and scepticism. *The Philosophical Review*, 103(1), 107–137.
- Gallois, A., & O’Leary-Hawthorne, J. (1996). Externalism and scepticism. *Philosophical Studies*, 81(1), 1–26.
- Gertler, B. (2015). Internalism, externalism and accessibilism. In S. C. Goldberg (Ed.), *Externalism, self-knowledge, and scepticism: New essays*. Cambridge University Press.
- Hawthorne, J. (2002). Deeply contingent a priori knowledge. *Philosophy and Phenomenological Research*, 65(2), 247–269.
- Jenkins, C. S., & Kasaki, M. (2015). The traditional conception of the a priori. *Synthese*, 192, 2725–2746.
- Kallestrup, J. (2011). Recent work on McKinsey’s paradox. *Analysis*, 71(1), 157–171.
- Kripke, S. (1980). *Naming and necessity*. Harvard University Press.
- McKinsey, M. (1991). Anti-individualism and privileged access. *Analysis*, 51(1), 9–16.
- McKinsey, M. (1994). Accepting the consequences of anti-individualism. *Analysis*, 54(2), 124–128.
- McKinsey, M. (2007). Externalism and privileged access are inconsistent. In B. McLaughlin & J. Cohen (Eds.), *Contemporary debates in philosophy of mind* (pp. 53–66). Blackwell.
- MgGinn, C. (1977). Charity, interpretation and belief. *Journal of Philosophy*, 74, 521–535.
- Pryor, J. (2007). What’s wrong with mckinsey-style reasoning? In S. Goldberg (Ed.), *Internalism and externalism in semantics and epistemology* (pp. 177–200). Oxford University Press.
- Putnam, H. (1973). Meaning and reference. *The Journal of Philosophy*, 70(19), 699–711.
- Putnam, H. (1975). The meaning of “Meaning.” In K. Gunderson (Ed.), *Language, mind, and knowledge*. (Vol. 2). Minnesota University Press.
- Putnam, H. (1981). *Reason truth and history*. Cambridge University Press.
- Ruffino, M. (2020). Contingent a priori truths and performatives. *Synthese*, 198(Suppl 22), S5593–S5613.
- Sawyer, S. (1998). Privileged access to the world. *Australasian Journal of Philosophy*, 76(4), 523–533.
- Sawyer, S. (2015). Contrastive self-knowledge and the McKinsey paradox. In S. C. Goldberg (Ed.), *Externalism, self knowledge, and scepticism* (pp. 75–94). Cambridge University Press.
- Thorpe, J. R., & Wright, C. J. (forthcoming). Putnam’s Proof Revisited. In J. Conant, & S. Chakraborty, *Engaging Putnam*. De Gruyter.
- Thorpe, J. R. (2019). Semantic self-knowledge and the vat argument. *Philosophical Studies*, 176(9), 2289–2306.
- Williamson, T. (1989). The contingent a priori: Has it anything to do with indexicals? *Analysis*, 46(3), 113–117.
- Wright, C. (2000). Cogency and question-begging: Some reflections on McKinsey’s paradox and Putnam’s proof. *Philosophical Issues*, 10, 140–163.
- Wright, C. (2011). McKinsey one more time. In A. Hatzimoysis (Ed.), *Self-knowledge* (pp. 80–104). Oxford University Press.